Adaptive TTL-Based Caching for Content Delivery

Soumya Basu[1], Aditya Sundarrajan[2], Javad Ghaderi[3], Sanjay Shakkottai[1], and Ramesh Sitaraman[2,4]

[1]The University of Texas at Austin

[2]University of Massachusetts Amherst

[3]Columbia University

[4]Akamai Technologies

arXiv:1704.04448v1 [cs.NI] 14 Apr 2017

ABSTRACT

Content Delivery Networks (CDNs) deliver a majority of the user-requested content on the Internet, including web pages, videos, and software downloads. A CDN server caches and serves the content requested by users. Designing caching algorithms that *automatically* adapt to the heterogeneity, burstiness, and non-stationary nature of real-world content requests is a major challenge and is the focus of our work. While there is much work on caching algorithms for stationary request traffic, the work on non-stationary request traffic is very limited. Consequently, most prior models are inaccurate for production CDN traffic that is non-stationary.

We propose two TTL-based caching algorithms and provide provable guarantees for content request traffic that is bursty and non-stationary. The first algorithm called d-TTL dynamically adapts a TTL parameter using a stochastic approximation approach. Given a feasible target hit rate, we show that the hit rate of d-TTL converges to its target value for a general class of bursty traffic that allows Markov dependence over time and non-stationary arrivals. The second algorithm called f-TTL uses two caches, each with its own TTL. The first-level cache adaptively filters out non-stationary traffic, while the second-level cache stores frequently-accessed stationary traffic. Given feasible targets for both the hit rate and the expected cache size, f-TTL asymptotically achieves both targets. We implement d-TTL and f-TTL and evaluate both algorithms using an extensive nine-day trace consisting of 500 million requests from a production CDN server. We show that both d-TTL and f-TTL converge to their hit rate targets with an error of about 1.3%. But, f-TTL requires a significantly smaller cache size than d-TTL to achieve the same hit rate, since it effectively filters out the non-stationary traffic for rarely-accessed objects.

1. INTRODUCTION

By caching and delivering content to millions of end users around the world, content delivery networks (CDNs) [1] are an integral part of the Internet infrastructure. A large CDN such as Akamai [2] serves several trillion user requests a day from 170,000+ servers located in 1500+ networks in 100+ countries around the world. The majority of today's Internet traffic is delivered by CDNs. CDNs are expected to deliver nearly two-thirds of the Internet traffic by 2020 [3].

The main function of a CDN server is to cache and serve content requested by users. The effectiveness of a caching algorithm is measured by its achieved hit rate in relation to its cache size. There are two primary ways of measuring the hit rate. The *object hit rate (OHR)* is the fraction of the requested objects that are served from cache and the *byte hit rate (BHR)* is the fraction of the requested content bytes that are served from cache. We devise algorithms capable of operating with both notions of hit rate in our work.

The major technical challenge in designing caching algorithms for a modern CDN is *adapting* to the sheer heterogeneity of the content that is accessed by users. The accessed content fall into multiple traffic classes that include web pages, videos, software downloads, interactive applications, and social networks. The classes differ widely in terms of the object size distributions and content access patterns. The popularity of the content also varies by several orders of magnitude with some objects accessed millions of times (e.g, an Apple iOS download), and other objects accessed one or twice (e.g, a photo in a Facebook gallery). In fact, as shown in Figure 2, 70% of the objects served by a CDN server are only requested once over a period of multiple days! Further, the requests served by a CDN server can change rapidly over time as different traffic mixes are routed to the server by the CDN's load balancer in response to Internet events.

Request statistics clearly play a key role in determining the hit rate of a CDN server. However, when request patterns vary rapidly across servers and time, a one-size-fits-all approach provides inferior hit rate performance in a production CDN setting. Further, manually tuning the caching algorithms for each individual server to account for the varying request statistics is prohibitively expensive. Thus, our goal is to devise self-tuning caching algorithms that can automatically learn and adapt to the request traffic and provably achieve any feasible hit rate and cache size, even when the request traffic is bursty and non-stationary.

Our work fulfills a long-standing deficiency in the current state-of-art in the modeling and analysis of caching algorithms. Even though real-world CDN traffic is known to be heterogeneous, with bursty (correlations over time) and non-stationary/transient request statistics, there are no known caching algorithms that provide theoretical performance guarantees for such traffic. In fact, much of the known formal models and analyses assume that the traffic follows the *Independent Reference Model* (IRM)[1]. Could it be that such models and analyses are sufficiently useful

[1]The inter arrival times are i.i.d. and the object request on each arrival are chosen independently from the same distribution.

in a production setting? We argue below with an example that this is not the case.

Deficiency of current models and analyses. A class of caching algorithms that has received much attention in the last decade are TTL-based caching algorithms [4, 5, 6, 7, 8, 9] that use a time-to-live (TTL) parameter to determine how long an object may remain in cache. Much is analytically known about the TTL caches when the traffic is assumed to follow IRM [6, 5]. The cornerstone of such analyses is Che's approximation [10] that provides a closed-form expression that relates the TTL value to the achieved hit rate and average cache size. Che's approximation is known to be accurate in cache simulations that use synthetic IRM traffic and is commonly used in the design of caching algorithms for that reason [11, 12, 8, 13, 14].

However, we show that Che's approximation produces erroneous results for actual production CDN traffic that is neither stationary nor IRM across the request. We used an extensive 9-day request trace from a production server in Akamai's CDN and derived TTL values for multiple hit rate and cache size targets using Che's approximation. We then simulated a cache with those TTL values on the production traces to derive the *actual* hit rate that was achieved. For a target hit rate of 60%, we observed that a fixed-TTL algorithm that uses the TTL computed from Che's approximation achieved a hit rate of 68.23% whereas the dynamic TTL algorithms proposed in this work achieve a hit rate of 59.36% (for a complete discussion, refer to Section 6.5). This difference between the target hit rate and that achieved by fixed-TTL highlights the inaccuracy of the current state-of-the-art theoretical modeling on production traffic. Similar observations have also been made in [15, 16], among others.

Our approach. We propose two TTL-based algorithms: d-TTL (for "dynamic TTL") and f-TTL (for "filtering TTL"). Rather than statically deriving the required TTL values by inferring the request statistics, our algorithms *dynamically* modify the TTL values after each request, incrementally adapting the TTLs to the current request patterns. Our algorithms do not have prior knowledge of request statistics; instead we adopt a stochastic approximation framework and ideas from actor-critic algorithms for parameter adaptation. To more accurately model CDN traces, we allow the request traffic to be non-independent and non-stationary. In particular, the request traffic can have Markov dependence over time, and further comprise of a mixture of stationary demands (statistics invariant over the timescale of interest), and non-stationary demands (finitely many requests, or in general requests with an asymptotically vanishing request rate). Further, we allow content to be classified into finitely many *types* of content, where each type has a target hit rate (OHR or BHR) and an average target cache size. In practice, a type could consist of all objects of a specific kind from a specific provider, e.g. CNN webpages, Facebook images, CNN video clips, etc.

Our Contributions. The main contributions follow.

1. d-TTL: A one-level TTL algorithm. Algorithm d-TTL maintains a single TTL value for each type, and dynamically adapts this value upon each arrival (new request) of an object of this type. Given a hit rate that is "feasible" (i.e. there exists a static genie-settable TTL parameter that can achieve this rate), we show that d-TTL almost surely converges to this given (feasible) hit rate target. Our result holds for a general class of bursty traffic (allowing Markov dependence over time), and even in the presence of non-stationary arrivals. To the best of our knowledge, this is the first adaptive TTL algorithm that can provably achieve a target hit rate with such stochastic traffic.

However, our empirical results with this algorithm shows that non-stationary and unpopular objects can contribute significantly to the cache size, while the hit rate gains with such objects residing in the cache remain negligible (heuristics employed in production CDNs that use Bloom filters to eliminate such traffic [17] support this observation).

2. f-TTL: A two-level TTL algorithm. The need to achieve a combined hit rate and cache size target motivates the f-TTL algorithm that comprises a pair of caches: a lower-level TTL cache that filters rare objects based on arrival history, and an higher-level adaptive TTL cache that stores filtered objects. We design an adaptation mechanism for a corresponding pair of TTL values (higher-level and lower-level) per type, and show that we can asymptotically achieve the desired hit rate (almost surely), under similar traffic conditions as with d-TTL (mixture of stationary and non-stationary traffic, with Markov dependence over time).

If the *stationary* part of the traffic is Poisson, we have the following stronger property. Given any feasible (hit rate, expected cache size) pair[2], the f-TTL algorithm asymptotically achieves a corresponding pair that dominates the given target[3]. Importantly, with non-stationary traffic, the two-level adaptive TTL strictly outperforms the one-level TTL cache with respect to the expected cache size.

Our proofs use a two-level stochastic approximation technique (along with a latent observer idea inspired from actor-critic algorithms [18]), and provides the first theoretical justification for the deployment of two-level caches such as ARC in production systems with non-stationary traffic.

3. Implementation and Empirical Evaluation: We implement both d-TTL and f-TTL algorithms and evaluate them using an extensive nine-day trace consisting of 500 million requests from a production Akamai CDN server. We observe that both d-TTL and f-TTL adapt well to the bursty and non-stationary production CDN request traffic. For a range of object hit rate targets, both d-TTL and f-TTL converge to that target with an error of about 1.3%. For a range of byte hit rate targets, both d-TTL and f-TTL converge to that target with an error that ranges from 0.3% to 2.3%. While the hit rate performance of both d-TTL and f-TTL are similar, f-TTL shows a distinct advantage in cache size due to its ability to filter out non-stationary traffic. In particular, f-TTL requires a cache that is 49% (resp., 39%) smaller than d-TTL to achieve the same object (resp., byte) hit rate. This renders f-TTL useful to CDN settings where large amounts of non-stationary traffic can be filtered out to conserve cache space while also achieving target hit rates.

Finally, from a practitioner's perspective, this work has the potential to enable new CDN pricing models. CDNs typically do not charge content providers on the basis of a guaranteed hit rate for their content, nor on the basis of the cache size that this used. Such pricing models have desirable properties, but do not commonly exist, in part, because current caching algorithms cannot provide such guarantees

[2]Feasibility here is with respect to any static two-level TTL algorithm of (target hit rate, expected cache size).

[3]A pair dominates another pair if hit rate is at least equal to the latter and expected size is at most equal to the latter.

with low overhead. Our caching algorithms are the first to provide a theoretical guarantee on hit rate for each content provider, while controlling the cache space that they can use. Thus, our work removes a technical impediment to hit rate and cache space based CDN pricing.

2. SYSTEM MODEL

2.1 Notation

In this paper, bold font characters indicate vector variables and normal font characters indicate scalar variables. We note $(x)^+ = \max(0, x)$, $\mathbb{N} = \{1, 2, \dots\}$, and $[n] = \{1, 2, \dots, n\}$. The equality among two vectors means component-wise equality holds. Similarly, inequality among two vectors (denoted by $\preccurlyeq$) means the inequality holds for each component separately.

2.2 Caching in CDNs

Every CDN server implements a cache that stores objects requested by users. When a user's request arrives at a CDN server, the requested object is served from its cache, if that object is present; otherwise, the CDN server fetches the object from a remote origin server that has original content and then serves it to the user. In addition, the CDN server may place the newly-fetched object in its cache. In general, a caching algorithm decides which object to place in cache, how long objects need to be stored in cache, and which objects should be evicted from cache.

When the requested object is found in cache, it is a *cache hit*, otherwise it is a *cache miss*. Additionally, it is often beneficial to maintain state (metadata such as the url of the object or an object ID) about a recently evicted object for some period of time. Then, we experience a *cache virtual hit* if the requested object is not in cache but its metadata is in cache. Note that metadata of an object takes much less cache space than the object itself.

The objective of a caching algorithm is to maximize the cache hit rate. On a cache hit, the object can be retrieved locally and served to the user quickly, whereas a cache miss would entail fetching the object from a remote origin server resulting in a slower response time. Further, having more cache hits reduces the content traffic (bits/sec) over the WAN from the content provider's origin to the CDN's server, resulting in cost savings for both the CDN and the content provider. Hit rate is an important metric for caches and there two commonly-used notions.

Definition 1. The ***object hit rate (OHR)*** is the ratio of the number of requests experiencing a *cache hit* to the total number of requests.

Definition 2. The ***byte hit rate (BHR)*** is the ratio of the number bytes experiencing a *cache hit* to the total bytes requested. Note that BHR directly measures the traffic reduction between the origin and the CDN.

The performance of a caching algorithm is best understood using its hit rate curve (HRC) that relates the hit rate that it achieves in relation to the cache size (in bytes) that it requires. A key metric that measures cache size is the ***size rate*** defined below.

Definition 3. The ***size rate*** of the system is the average size of the cache per unit arrival rate.

To illustrate the definition, consider the size rate of 10 sec and the arrival rate of 10 Gbps. This implies that the average cache size is 100 Gb. A major objective of this paper is to design a caching algorithm this is oblivious to the arrival process and achieves a target object (or, byte) hit rate. Another related and more ambitious goal is to achieve both a given target hit rate and target size rate.

2.3 Content Request Model

There are different types of content served by modern CDNs. A content type may represent a specific genre of content (videos, web pages, etc.) from a specific content provider (CNN, Facebook, etc.). As such, a single server cache could be shared among dozens of content types. A salient feature of traffic in CDNs is that the objects belonging to one type can be very different from objects belonging to another type, in terms of their popularity characteristics, request patterns and object size distributions. Most content types exhibit a long tail of popularity where there is a smaller set of *recurring objects* that demonstrate a stationary behavior in their popularity and are requested frequently by users, and a larger set of *rare objects* that are unpopular and show a high degree of non-stationarity. Examples of rare objects include those that are requested infrequently or even just once, a.k.a. one-hit wonders [19]. Another example is an object that is rare in a temporal sense and is frequently accessed within a small time window, but is seldom accessed again. Such a bursty request pattern can occur during flash crowds [20]. In this section, we present a content request model that captures these characteristics. We make the following assumptions in our content request model.

Assumption 1. Description of Content:

- There are T types of content, where T is finite.
- The universe of each type $t \in T$ is represented as $\mathcal{U}_t$ and consists of both recurring objects and rare objects. Let the set of recurring objects be $\mathcal{K}_t$ with $K_t (< \infty)$ different objects. The remaining objects are rare objects which are (potentially) countably infinite in number.
- The entire universe of objects is represented as $\mathcal{U} \equiv \cup_{t \in T} \mathcal{U}_t$, whereas the recurring objects are represented as a finite set $\mathcal{K} \equiv \cup_{t \in T} \mathcal{K}_t$. Let $K \equiv |\mathcal{K}| = \sum_{t \in T} K_t$.
- We adopt a unifying nomenclature for all the objects, where each object $c \in \mathcal{U}$ is represented by a tuple, $c = (c_i, c_t, c_m)$, where c_i is the unique *label* for the object (e.g., its url), c_t is the *type* that the object belongs to, and c_m is the actual body of the object c. Note that c_i and c_t constitute the metadata of object c which requires very small cache space, while the bulk of cache space is used to store c_m. We denote the size of object c as $w_c = |c_m|$, where $w_c \leq w_{\max}$ for all c. Unless specified otherwise, caching object c means caching the tuple (c_i, c_t, c_m). We differentiate between object $c = (c_i, c_t, c_m)$ and object metadata $\tilde{c} = (c_i, c_t)$. Further note that the object metadata can be fully extracted from the incoming request. For simplicity, we assume all rare objects of type t have equal size $\bar{w}_t$.[4]

[4]This could be relaxed to *average* size for type t rare objects, as long as the average size over a large enough time window has $o(1)$ difference from the average, w.p. 1.

2.3.1 General Content Request Model

The content request model has two components: (1) the arrival process, and (2) the object labeling process. The requests arrive following a continuous-time arrival process, and upon each arrival, the label of the requested object is generated by the object labeling process. The arrival process is assumed to have i.i.d. inter-arrival times, however, the *object labeling process has a non-stationary component* as the *rare object* labels can be chosen arbitrarily. Further, there could be a Markovian dependence on the labels of subsequent object requests. Formally, our content request model is built on a *Markov renewal process* [21] (to model the stationary components) followed by rare object labeling to model non-stationary components, as described below.

Assumption 2. General Content Request Model:

- **A Markov renewal process.** Let $A(l)$ be the arrival time of the l-th request. Let $Z(l)$ be the object label, if the l-th arrival is a recurrent object. Otherwise, if the l-th arrival is a rare object, let $Z(l)$ be its object type. The l-th inter-arrival time is defined by $X(l) = A(l) - A(l-1)$. We assume that the process $(A(l), Z(l))_{l\in\mathbb{N}}$ evolves as a Markov renewal process [21]. The inter-arrival time process $X(l)$, the process $Z(l)$, and the actual labeling of recurrent and rare objects are described below.

- **Inter-arrival process $X(l)$.** The inter-arrival times $(X(l))_{l\in\mathbb{N}}$ are identically distributed, independently of each other and $Z(l)$. The inter-arrival time distribution follows a probability density function (p.d.f.), $f(x)$ which is *absolutely continuous* w.r.t a Lebesgue measure on $(\mathbb{R}, +)$ and has *simply connected support*. The inter-arrival time has a *nonzero finite mean* denoted by $1/\lambda$.

- **The process $Z(l)$.** The process $Z(l)$ is a Markov chain over $(K+T)$ states, indexed by $1, \cdots, K+T$. The first K states represent the K recurring objects. The rare objects (possibly infinite in number) are grouped according to their types, thus producing the remaining T states, i.e, the states $K+1, \cdots, K+T$ represent rare objects of types $1, \cdots, T$, respectively. The transition probability matrix of the Markov chain $Z(l)$ is given by P, where

  $$P(c, c') := P(Z(l) = c' | Z(l-1) = c),\ \forall c, c' \in [K+T].$$

  We assume that the diagonal entries $P(c, c) > 0$. Also the Markov chain is assumed to be irreducible and aperiodic, thus possesses a stationary distribution denoted by $\boldsymbol{\pi}$.

- **Object labeling process.**

  Recurrent objects: On the l-th arrival, if the Markov chain $Z(l)$ moves to a state $k \in [K]$, the arrival is labeled by the recurrent object k.

  Rare objects: On the l-th arrival, if the Markov chain $Z(l)$ moves to a state $K+t$, $t \in [T]$, the arrival is labeled by a rare object of type t, chosen from the group of rare objects of type t, such that the label assignment has no stationary behavior in the time-scale of $O(1)$ arrivals and it shows a *rarity* behavior in large time-scales.

  Formally, given a constant $R > 0$, and type $t \in [T]$, define $\beta_t(l; R)$ as the indicator function of the event that: (1) the l-th arrival is any particular rare object c of type t, *and* (2) the previous arrival of the same rare object c happened (strictly) less than R units of time earlier. Note that $\beta_t(l; R)$ does not depend on a specific $c \in t$; the definition of $\beta_t(l; R)$ implies that it accumulates over all contents in type t (i.e. it is the summation of indicators over all $c \in t$). The labeling of rare objects has to respect the following *R-rarity condition*.

  Definition 4 (R-rarity condition)**.** For any type $t \in T$, and a finite $R > 0$,

  $$\lim_{m\to\infty} \frac{1}{N_m^t} \sum_{l=m}^{m+N_m^t} \beta_t(l; R) = 0,\ \text{w.p. 1.} \tag{1}$$

  for any $N_m^t = \omega(\sqrt{m})$.

In words, the "R-rarity condition" states that burst requests for a rare object can arrive over time (bursts are separated apart by at least R time units by definition), but they should become asymptotically less frequent as time goes to infinity. Note that the "R-rarity condition" is not specified for a particular constant R, but it is parameterized by R and we shall specify the specific value later. If R-rarity condition holds then R'-rarity condition also holds for any $R' \in [0, R)$, which easily follows from the definition. Also, for any type t, in the long run, the aggregate fraction of the total arrivals which are rare objects of type t is $\alpha_t := \pi_{(K+t)}$, under the stationary distribution π of the $Z(l)$ process. If $\alpha_t > 0$, then for the rarity condition to hold, there must be infinitely many rare contents of the same type t (over an infinite time horizon).

Remark 1 (Relevance of the content request model). We note that most of the popular inter-arrival distributions, e.g., Exponential, Phase-type, Weibull, satisfy the inter-arrival model in Assumption 2. Moreover, it is easy to see that any i.i.d. distribution for content popularity, including Zipfian distribution, is a special case of our object labeling process. In fact, the labeling process is much more general in the sense that it can capture the influence of different objects on other objects, which may span across various types.

Remark 2 (Relevance of the R-rarity condition). The proposed rarity condition is fairly general, as at any point in time, no matter how large, it allows the existence of rare objects which may exhibit burstiness for small time windows. In particular, the following real-world scenarios satisfy the rarity condition (1):

- *One-hit wonders [19].* Suppose a constant fraction α of total arrivals consists of rare objects that are requested only once. This clearly satisfies (1) as the indicator $\beta_t(l; R)$ is zero for the first and only time that an object is requested.

- *Flash crowds [20].* Constant size bursts of requests for rare objects may occur over time, with $O(\sqrt{\tau})$ number of such bursts up to time τ. This allows for infinitely many such bursts. In this scenario, almost surely, for any window of length N, starting from a time τ_0, the total number of requests for rare objects inside the window is less than $O(\sqrt{\tau_0 + N})$, which is $o(N)$ for $N = \omega(\sqrt{\tau_0})$. Therefore, (1) is satisfied.

Further, we can weaken the rarity condition by requiring the condition to hold with high probability (w.r.t. N_m^t, for each t) instead of w.p. 1.

2.3.2 Poisson Arrival with Independent Labeling

We next consider a specific model for the arrival process which is an interesting and well studied special case of Assumption 2. We will later show under this arrival process we can achieve stronger guarantees on the system performance.

Assumption 3. Poisson Arrival with Independent Labeling:

- The inter arrival times are i.i.d. and follows *e*xponential distribution with rate $\lambda > 0$.
- The labels for the recurring objects are given independently. For each recurrent object c, the probability of assigning label c equals π_c and its size is given as w_c, which is non-decreasing w.r.t. probability π_c. Whereas, on each turn, some rare object of type t is chosen with probability α_t and following the condition in (1). For each type $t \in [T]$, all rare objects of type t have size $\bar{w}_t$.

3. TTL-BASED CACHING ALGORITHMS

A TTL-based caching algorithm works as follows. When a new object is requested, the object is placed in cache and is associated with a time-to-live (TTL) value. If no new requests are received for that object, the TTL value is decremented in real-time and the object is evicted when the TTL becomes zero. If a cached object is requested, the TTL is reset to its original value.

The fundamental challenge in cache design is striking a balance between two conflicting objectives: minimizing the cache size and maximizing the cache hit rate. In a TTL cache, the TTL helps balance the cache size and hit rate objectives. When the TTL increases, the object stays in cache for a longer period of time, increasing the cache hit rate, at the expense of a larger cache size. The opposite happens when TTL decreases.

In addition, it is desirable to allow different Quality of Service (QoS) guarantees for different types of objects, i.e., different *cache hit rates* and *size targets* for different types of objects. For example, a lower hit rate that results in a higher response time may be tolerable for a software download that happens in the background. But, a higher hit rate that results in a faster response time is desirable for a web page that is delivered in real-time to the user. As seen in previous work [4, 13], with full knowledge of a stationary arrival process, it is possible to specify the TTL to meet the target hit rate. However, without the prior knowledge of the arrival process, the desired TTL has to be learned from the request arrivals, while also providing the QoS guarantees. In this work, we show how we can tune the TTL value to asymptotically achieve a *target hit rate vector* $\mathbf{h}^*$ and a (feasible) *target size rate vector* $\boldsymbol{s}^*$. The t-th components of $\mathbf{h}^*$ and $\boldsymbol{s}^*$, i.e., h_t^* and s_t^* respectively, denote the target hit rate and the target size rate for objects of type $t \in [T]$. Note that the CDN can group objects into types in an arbitrary way. But, if the objective is to achieve an overall hit rate and cache size, all objects can be grouped into a single type. It should also be noted that the algorithms proposed in this work *do not* try to achieve the target hit rate with the smallest cache size; this is a non-convex optimization problem that is not the focus of this work. Instead, we only try to achieve a *given* target hit rate and target size rate by learning the parameters of the caching system.

We propose two adaptive TTL algorithms. First, we present a dynamic TTL algorithm (d-TTL) that adapts its TTL to achieve a target hit rate $\mathbf{h}^*$. d-TTL dynamically increases the TTL when the current hit rate is below the target $\mathbf{h}^*$, and decreases the TTL when the current hit rate is above the target $\mathbf{h}^*$, using stochastic approximation techniques.

While d-TTL does a good job of achieving the target hit rate, it does this at the expense of caching rare and unpopular recurring content for an extended period of time, thus causing an increase in cache size without any significant contribution towards the cache hit rate. We present a second adaptive TTL algorithm called filtering TTL (f-TTL) that filters out rare content to achieve the target hit rate with a smaller cache size. f-TTL is a two-level TTL algorithm that filters out rare and unpopular content to achieve both a given target hit rate $\mathbf{h}^*$ and a feasible target size rate $\mathbf{s}^*$.

To the best of our knowledge, both d-TTL and f-TTL are the first adaptive TTL-based caching algorithms that are able to achieve a target hit rate and a feasible target size rate for non-stationary traffic. We discuss d-TTL and f-TTL in detail in the following sections.

3.1 Dynamic TTL (d-TTL)

We propose a dynamic TTL (d-TTL) algorithm where at each arrival, the TTL is updated using stochastic approximation to achieve a target hit rate $\mathbf{h}^*$. The d-TTL algorithm converges to the target hit rate by increasing the TTL during cache misses and decreasing the TTL during cache hits.

Let every object c present in the cache, $\mathcal{C}$, have a timer ψ_c^0 that encodes its remaining TTL. We denote by $\boldsymbol{\theta}(l) \in \mathrm{R}^T$, the TTL at the l-th arrival. Note that $\theta_t(\cdot)$ represents the TTL for objects of type t. On the l-th arrival, if the requested object c of type t is present in cache, $\theta_t(l)$ is decremented, and if the requested object c to type t is not present in cache, object c is fetched from the origin, cached in the server and $\theta_t(l)$ is incremented. In both cases, ψ_c^0 is set to the updated timer $\theta_t(l+1)$ until the object is re-requested or evicted.

We restrict the TTL value $\boldsymbol{\theta}(l)$ to $\boldsymbol{\theta}(l) \preceq \boldsymbol{L}$[5]. Here $\mathbf{L}$ is the truncation parameter of the algorithm and an increase in $\mathbf{L}$ increases the achievable hit rate (see Section 4 for details). For the convenience of exposition, we introduce a latent variable $\boldsymbol{\vartheta}(l) \in \mathbb{R}^T$ where $\boldsymbol{\vartheta}(l) \in [0,1]$. Without loss of generality, instead of adapting $\boldsymbol{\theta}(l)$, we dynamically adapt each component of $\boldsymbol{\vartheta}(l)$ and set $\theta_t(l) = L_t \vartheta_t(l)$, where $\vartheta_t(\cdot)$ is the latent variable for objects of type t. The d-TTL algorithm for dynamically changing the value of $\boldsymbol{\theta}(l)$ is presented in Algorithm 1. As previously discussed, object c of type t is evicted from the cache when $\psi_c^0 = 0$.

3.2 Filtering TTL (f-TTL)

The d-TTL algorithm achieves the target hit rate at the expense of caching rare and unpopular content. We propose a two-level filtering TTL algorithm (f-TTL) that filters non-stationary content to achieve the target hit rate along with the target size rate. We briefly discuss the idea behind the two-level algorithm before presenting the complete algorithm.

The two-level f-TTL algorithm maintains two caches, a lower-level cache $\mathcal{C}_s$ and a higher-level cache $\mathcal{C}$. To facili-

[5]This gives an upper bound, typically a large one, over the size of the cache. Further it can control the staleness of objects.

Algorithm 1 Dynamic TTL (d-TTL)

Input:
Target hit rate $\mathbf{h}^*$, TTL upper bound $\mathbf{L}$.
For l-th request, $l \in \mathbb{N}$, object label $c(l)$, size $w(l)$ & type $t(l)$.
Output: Cache requested object using dynamic TTL, $\boldsymbol{\theta}$.

Initialize: Latent variable $\boldsymbol{\vartheta}(0) = \mathbf{0}$.
for all $l \in \mathbb{N}$ **do**
  if Cache hit, $c(l) \in \mathcal{C}$ **then**
    $Y(l) = 1$
  else Cache miss
    $Y(l) = 0$
  Update TTL $\theta_{t(l)}(l)$:

$$\vartheta_{t(l)}(l+1) = \mathcal{P}_{[0,1]}\left(\vartheta_{t(l)}(l) + \eta(l)\widehat{w}(l)\left(h^*_{t(l)} - Y(l)\right)\right)$$
$$\theta_{t(l)}(l+1) = L_{t(l)}\vartheta_{t(l)}(l+1) \tag{2}$$

  where,
    $\eta(l) = \frac{\eta_0}{l^\alpha}$ is a decaying step size for $\alpha \in (1/2, 1)$,
    $\mathcal{P}_{[0,1]}(x) = \min\{1, \max\{0, x\}\}$,
    $\widehat{w}(l) = 1$ for OHR and $w(l)$ for BHR.

  Cache c with TTL $\psi^0_{c(l)} = \theta_{t(l)}(l+1)$ in $\mathcal{C}$.

tate the filtering process, it uses two dynamically adapted TTL values—one smaller or equal to the other. Upon a *cache miss* for object c, the object c, potentially unpopular, is cached with the smaller TTL to ensure quick eviction while the metadata $\tilde{c}$ is cached with the larger TTL to retain *memory* of this request. This *memory* is akin to shadow caches used previously (e.g. ARC [22]). Whereas, the quick eviction of newly arrived object effectively filters out rare and unpopular content, similar to Bloom filtering [19] or 2Q [23]. A *cache hit*, as well as, a *cache virtual hit* of a requested object c, ascertain its popularity. In both cases, object c is cached directly in $\mathcal{C}$ with the larger TTL. The primary cache $\mathcal{C}$ behaves similar to the single level cache in d-TTL (Algorithm 1), with the exception that it stores mostly stationary content. Thus, f-TTL utilizes the primary cache more efficiently to simultaneously achieve a target hit rate, and a target size rate.

TTL timers for f-TTL. Every object in f-TTL is associated with a TTL tuple $(\psi^0_c, \psi^1_c, \psi^2_c)$, where ψ^0_c is the TTL associated with the object cached in $\mathcal{C}$, ψ^1_c is the TTL associated with the object cached in $\mathcal{C}_s$ and ψ^2_c is the TTL associated with the metadata of the object cached in $\mathcal{C}_s$. We maintain the invariant that $\psi^1_c \le \psi^2_c$. At any point in time, the TTL tuple is either $(\psi^0_c, 0, 0)$—object $c \in \mathcal{C}$, or $(0, 0, \psi^2_c)$—object $c \notin \mathcal{C} \cup \mathcal{C}_s$ and metadata $\tilde{c} \in \mathcal{C}_s$, or $(0, \psi^1_c, \psi^2_c)$—object $c \in \mathcal{C}_s$ and metadata $\tilde{c} \in \mathcal{C}_s$. To capture the dynamics of $\mathcal{C}_s$, f-TTL maintains an additional time varying TTL $\boldsymbol{\theta}^s(l)$ along with $\boldsymbol{\theta}(l)$, to achieve the target size rate. The TTL timers are dynamically updated as follows.

Consider the l-th arrival and suppose the request is for object c (of type t). Then, during a cache hit we cache object c in cache $\mathcal{C}$ with TTL tuple $(\theta_t(l), 0, 0)$. If the hit occurs when c is in $\mathcal{C}_s$, we delete c and its metadata $\tilde{c}$ from $\mathcal{C}_s$. During a cache virtual hit, we cache c in $\mathcal{C}$ with TTL tuple $(\theta_t(l), 0, 0)$ and delete $\tilde{c}$ from $\mathcal{C}_s$. Finally, during a cache miss, we cache object c in $\mathcal{C}_s$ with TTL tuple $(0, \theta^s_t(l), \theta_t(l))$, where the object c and its meta data $\tilde{c}$ are cached in $\mathcal{C}_s$ with timer $\theta^s_t(l)$ and $\theta_t(l)$, respectively. Object c is evicted from $\mathcal{C}$ (resp., $\mathcal{C}_s$) when ψ^0_c (resp., ψ^1_c) becomes 0. Further, the metadata $\tilde{c}$ is evicted from $\mathcal{C}_s$ when ψ^2_c is 0. See Algorithm 2 for a detailed description of f-TTL.

Algorithm 2 Filtering TTL (f-TTL)

Input:
Target hit rate $\mathbf{h}^*$, target size rate $\mathbf{s}^*$, TTL upper bound $\mathbf{L}$.
For l-th request, $l \in \mathbb{N}$, object label $c(l)$, size $w(l)$ & type $t(l)$.
Output: Cache requested object using dynamic TTLs, $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^s$.

Initialize: Latent variables, $\boldsymbol{\vartheta}(0) = \boldsymbol{\vartheta}^s(0) = \mathbf{0}$.
for all $l \in \mathbb{N}$ **do**
  if Cache hit, $c(l) \in \mathcal{C} \cup \mathcal{C}_s$ **then**
    $Y(l) = 1$,
    $s(l) = \begin{cases} \theta_{t(l)}(l) - \psi^0_{c(l)}, & \text{if } c \in \mathcal{C} \\ \theta_{t(l)}(l) - \psi^1_{c(l)}, & \text{if } c \in \mathcal{C}_s. \end{cases}$
  else if Virtual hit, $c(l) \notin \mathcal{C} \cup \mathcal{C}_s$ and $\tilde{c}(l) \in \mathcal{C}_s$ **then**
    $Y(l) = 0$, $s(l) = \theta_{t(l)}(l)$.
  else Cache miss
    $Y(l) = 0$, $s(l) = \theta^s_{t(l)}(l)$.
  Update TTL $\theta_{t(l)}(l)$:

$$\vartheta_{t(l)}(l+1) = \mathcal{P}_{[0,1]}\left(\vartheta_{t(l)}(l) + \eta(l)\widehat{w}(l)\left(h^*_{t(l)} - Y(l)\right)\right)$$
$$\theta_{t(l)}(l+1) = L_{t(l)}\vartheta_{t(l)}(l+1),$$

  where,
    $\eta(l) = \frac{\eta_0}{l^\alpha}$ is a decaying step size for $\alpha \in (1/2, 1)$,
    $\mathcal{P}_{[0,1]}(x) = \min\{1, \max\{0, x\}\}$,
    $\widehat{w}(l) = 1$ for OHR and $w(l)$ for BHR.

  Update TTL $\theta^s_{t(l)}(l)$:

$$\vartheta^s_{t(l)}(l+1) = \mathcal{P}_{[0,1]}\left(\vartheta^s_{t(l)}(l) + \eta_s(l)w(l)(s^*_{t(l)} - s(l))\right)$$
$$\theta^s_{t(l)}(l+1) = L_{t(l)}\vartheta_{t(l)}(l+1)\Gamma\left(\vartheta_{t(l)}(l+1), \vartheta^s_{t(l)}(l+1); \epsilon\right)$$

  where,
    $\eta_s(l) = \frac{\eta_0}{l}$ and ϵ is a parameter of the algorithm,
    $\Gamma(\cdot, \cdot; \epsilon)$ is a *threshold function*.

  if Cache hit, $c(l) \in \mathcal{C} \cup \mathcal{C}_s$ **then**
    if $c(l) \in \mathcal{C}_s$ **then**
      Evict $\tilde{c}(l)$ from $\mathcal{C}_s$ and move $c(l)$ from $\mathcal{C}_s$ to $\mathcal{C}$.
    Set TTL tuple to $(\theta_{t(l)}(l+1), 0, 0)$.
  else if Virtual hit, $c(l) \notin \mathcal{C} \cup \mathcal{C}_s$ and $\tilde{c}(l) \in \mathcal{C}_s$ **then**
    Evict $\tilde{c}(l)$ from $\mathcal{C}_s$,
    Cache $c(l)$ in $\mathcal{C}$ and set TTL tuple to $(\theta_{t(l)}(l+1), 0, 0)$.
  else Cache Miss
    Cache $c(l)$ and $\tilde{c}(l)$ in $\mathcal{C}_s$ and
    set TTL tuple to $(0, \theta^s_{t(l)}(l+1), \theta_{t(l)}(l+1))$.

Adapting $\boldsymbol{\theta}^s(l)$ and $\boldsymbol{\theta}(l)$. The update rules for the parameters $\boldsymbol{\vartheta}(l)$ and $\boldsymbol{\theta}(l)$ in f-TTL are the same as those in d-TTL, both of which try to achieve the target hit rate. However, the update rule for $\boldsymbol{\theta}^s(l)$ depends on $\boldsymbol{\vartheta}(l)$ and $\boldsymbol{\vartheta}^s(l)$, where the latter is a new latent variable that lies in the range $[0, 1]^T$, and is dynamically updated to achieve the target size rate $\boldsymbol{s}^*$. Before describing this update rule, we first define a twice differentiable non-decreasing *threshold function*, $\Gamma(x, y; \epsilon) : [0, 1]^2 \to [0, 1]$, which takes value 1 for $x \ge 1 - \epsilon/2$ and takes value y for $x \le 1 - 3\epsilon/2$, and its slope in between is bounded by $4/\epsilon$. One such function can be given as follows with the convention $0/0 = 1$,

$$\Gamma(x, y; \epsilon) = \left(y + \frac{(1-y)((x - 1 + \frac{3\epsilon}{2})^+)^4}{((x - 1 + \frac{3\epsilon}{2})^+)^4 + ((1 - \frac{\epsilon}{2} - x)^+)^4}\right).$$

The parameter $\boldsymbol{\theta}^s(l)$ is defined as

$$\theta^s_t(l) = L_t\vartheta_t(l)\Gamma\left(\vartheta_t(l), \vartheta^s_t(l); \epsilon\right) \ \forall t \in [T].$$

This definition maintains the invariant $\theta^s_t(l) \le \theta_t(l)$ $\forall t, l$.

Note that, in the extreme case when $\boldsymbol{\vartheta}^s(l) = 0$, we only cache the metadata of the requested object on first access, but not the object itself. We call this the *full filtering TTL*.

Like d-TTL, the f-TTL algorithm adaptively decreases $\boldsymbol{\theta}(l)$ during cache hits and increases $\boldsymbol{\theta}(l)$ during cache misses. Additionally, f-TTL also increases $\boldsymbol{\theta}(l)$ during cache virtual hits. This ensures that stationary objects in cache $\mathcal{C}_s$ are cached for the longer duration $\boldsymbol{\theta}(l)$ in $\mathcal{C}$. In addition, f-TTL has a secondary target of achieving a target size rate through proper filtering which is accomplished by adapting $\boldsymbol{\theta}^s(l)$. The update rule for $\boldsymbol{\theta}^s(l)$ is given in Equation 4 in Algorithm 2 through the latent variable $\boldsymbol{\vartheta}^s(l)$. The complete algorithm for f-TTL is given in Algorithm 2.

Remark 3 (The role of the threshold function $\Gamma(x, y; \epsilon)$). Given a target hit rate, target size rates that are very small can be infeasible. In the absence of the threshold function $\Gamma(\vartheta_{t(l)}, \vartheta^s_{t(l)}; \epsilon)$, when the target size rate is too small, f-TTL moves to the full filtering mode (i.e., $\theta^s_t = 0$) which may prevent us from achieving hit rates that are feasible by the d-TTL algorithm. In this case, the threshold function disregards the size dynamics and increases θ^s_t gradually to ensure that the target hit rate is achieved. The latent variable $\boldsymbol{\vartheta}^s$ ensures that the size rate estimates are unbiased while achieving the target hit rate. The use of such update rules is reminiscent of actor-critic algorithms in reinforcement learning [18].

4. ANALYSIS OF TTL-BASED CACHING

In this section we present our main theoretical results. We consider a setting where the TTL parameters live in a compact space. Indeed, if the TTL values become unbounded, then objects never leave the cache after entering it. This setting is captured through the definition of $\mathbf{L}$ feasibility, presented below.

Definition 5. For an arrival process $\mathcal{A}$ and d-TTL algorithm, object (byte) hit rate $\mathbf{h}$ is '$\mathbf{L}$-feasible' if there exists a $\boldsymbol{\theta} \preccurlyeq \mathbf{L}$ such that d-TTL algorithm with fixed TTL $\boldsymbol{\theta}$ achieves $\mathbf{h}$ asymptotically almost surely under $\mathcal{A}$.

Definition 6. For an arrival process $\mathcal{A}$ and f-TTL caching algorithm, object (byte) hit rate, size rate tuple $(\boldsymbol{h}, \boldsymbol{s})$ is '$\mathbf{L}$-feasible' if there exist $\boldsymbol{\theta}_1 \preccurlyeq \mathbf{L}$ and $\boldsymbol{\theta}_2 \preccurlyeq \mathbf{L}$, such that f-TTL algorithm with fixed TTL pair $(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ achieves $(\boldsymbol{h}, \boldsymbol{s})$ asymptotically almost surely under $\mathcal{A}$.

To avoid trivial cases (hit rate being 0 or 1), we have the following definition.

Definition 7. A hit rate $\mathbf{h}$ is 'typical' if $h_t \in (0, 1)$ for all types $t \in [T]$.

4.1 Main Results

We now show that both d-TTL and f-TLL asymptotically almost surely (a.a.s.) achieve any 'feasible' *object (byte) hit rate*, $\mathbf{h}^*$ for the arrival process in Assumption 2, using stochastic approximation techniques. Further, we prove a.a.s that f-TTL converges to a specified $(\boldsymbol{h}^*, \boldsymbol{s}^*)$ tuple for *object (byte) hit rate* and *size rate*.

Theorem 1. ***d-TTL:*** *Under Assumption 2 with $\|\boldsymbol{L}\|_\infty$-rarity condition, if hit rate target $\boldsymbol{h}^*$ is both $\boldsymbol{L}$-feasible and 'typical', the d-TTL algorithm with parameter $\boldsymbol{L}$ converges to a TTL value of $\boldsymbol{\theta}^*$ a.a.s. Further, the average hit rate converges to $\boldsymbol{h}^*$ a.a.s.*

f-TTL: *Suppose that the target tuple of hit rate and size rate, $(\boldsymbol{h}^*, \boldsymbol{s}^*)$, is $(1-2\epsilon)\boldsymbol{L}$-feasible with $\epsilon > 0$ and additionally, $\boldsymbol{h}^*$ is 'typical'. Under Assumption 2 with $\|\boldsymbol{L}\|_\infty$-rarity condition, the f-TTL algorithm with parameter $\boldsymbol{L}$ and ϵ converges to a TTL pair $(\boldsymbol{\theta}^*, \boldsymbol{\theta}^{s^*})$ a.a.s. Further the average hit rate converges to $\boldsymbol{h}^*$ a.a.s., while the average size rate converges to some $\hat{\boldsymbol{s}}$ a.a.s.. Additionally, $\hat{\boldsymbol{s}}$ for each type t, satisfies one of the following three conditions:*

1. *The average size rate converges to $\hat{s}_t = s^*_t$ a.a.s.*
2. *The average size rate converges to $\hat{s}_t > s^*_t$ a.a.s. and $\theta^{s^*}_t = 0$ a.a.s.*
3. *The average size rate converges to $\hat{s}_t < s^*_t$ a.a.s. and $\theta^{s^*}_t = \theta^*_t$ a.a.s.*

Remark 4. We call the second scenario *collapse to full filtering TTL*, because in this case, the lower-level cache contains only labels of objects instead of caching the objects themselves. We refer to the third scenario as *collapse to d-TTL*, because in this case, cached objects have equal TTL values in both the higher and lower-level caches.

Proof Sketch. Here we present a proof sketch of Theorem 1. The complete proof can be found in Appendix A. The proof of Theorem 1 consists of two parts. The first part deals with the 'static analysis' of the caching process, where parameters $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$ both take fixed values in $[0, 1]$ (i.e., no adaptation of parameters). In the second part (the 'dynamic analysis'), employing techniques from the theory of stochastic approximation [24], we show that the TTL $\boldsymbol{\theta}$ for d-TTL and the TTL pair $(\boldsymbol{\theta}, \boldsymbol{\theta}^s)$ for f-TTL converges almost surely. Further, the average hit rate (and average size rate for f-TTL) satisfies Theorem 1.

The evolution of the caching process is represented as a discrete time stochastic process uniformized over the arrivals into the system. At each arrival, the system state is completely described by the following: *(1)* the timers of recurrent objects (i.e. $(\psi^0_c, \psi^1_c, \psi^2_c)$ for $c \in \mathcal{K}$), *(2)* the current value of the pair $(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s)$, and *(3)* the object requested on the last arrival. However, due to the presence of a constant fraction of non-stationary arrivals in Assumption 2, we maintain a state with incomplete information. Specifically, our (incomplete) state representation does not contain the timer values of the rare objects present in the system. This introduces a bias (which is treated as noise) between the actual process, and the evolution of the system under incomplete state information.

In the static analysis, we prove that the system with fixed $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$ exhibits uniform convergence to a unique stationary distribution. Further, using techniques from regeneration process and the 'rarity condition' (1), we calculate the asymptotic average hit probabilities and the asymptotic average size rates of each type for the original caching process. We then argue that asymptotic averages for both the hit rate and size rate of the incomplete state system is same as the original system. This is important for the dynamic analysis because this characterizes the time averages of the adaptation of $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$.

In the dynamic analysis, we analyze the system under variable $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$, using results of almost sure convergence of (actor-critic) stochastic approximations with a two timescale separation [25]. The proof follows the ODE method; the following are the key steps in the proof of dynamic analysis:

1. We show that the effects of the bias introduced by the non-stationary process satisfies Kushner-Clark condition [24]

2. The average (w.r.t. the history up to step l) of the l-th update as a function of $(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s)$ is Lipschitz continuous

3. The incomplete information system is uniformly ergodic

4. The ODE (depends on $\boldsymbol{\vartheta}^s$) representing the mean evolution of $\boldsymbol{\vartheta}$ has a unique stationary point which can be represented as a unique function of $\boldsymbol{\vartheta}^s$

5. (f-TTL analysis with two timescales) Let the ODE at the slower time scale, representing the mean evolution of $\boldsymbol{\vartheta}^s$, have stationary points $\{(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s)_i\}$. We characterize each stationary point, and show that it corresponds to one of the three cases stated in Theorem 1.

Finally, we remark on a couple of ideas that are key in formally proving the results. First, varying $\boldsymbol{\vartheta}$ at a faster timescale compared to $\boldsymbol{\vartheta^s}$ enables us to achieve a targeted hit rate (a.s.) (this is the main goal) for both f-TTL and d-TTL. Second, the proof utilizes ideas from actor-critic algorithms [18], where the parameters labeled as 'actors' dictate the policy of the algorithm, whereas the parameters labeled as 'critics' learn and critique the current policy. In the f-TTL algorithm, the TTL parameters, $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^s$, play the role of actors by controlling the operation of the higher and lower-level caches. The latent parameters $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$, are the critics that learn the empirical averages of hit rate and size rate, respectively.

As stated in Theorem 1, the f-TTL algorithm converges to one of three scenarios. We note that scenario 1 and scenario 3 have 'good' properties. Specifically, in each of these two cases, the f-TTL algorithm converges to an average size rate which is smaller or equal to the target size rate. However, in scenario 2, the average size rate converges to a size rate larger than the given target under general arrivals in Assumption 2. However, under Assumption 3, we show that the scenario 2 cannot occur, as formalized in Lemma 1 and its Corollary.

Lemma 1. *Under Assumption 3 with L-rarity condition and for any type t, suppose f-TTL algorithm achieves an average hit rate h_t with two different TTL pairs, 1) (θ_t, θ_t^s) with $\theta_t^s = 0$ (full filtering), and 2) $(\hat{\theta}_t, \hat{\theta}_t^s)$, with $\hat{\theta}_t^s > 0$, where $\max\{\theta_t, \hat{\theta}_t\} \leq L$. Then the average size rate achieved with the first pair is less or equal to the average size rate achieved with the second pair. Moreover, in the presence of rare objects of type t, i.e. $\alpha_t > 0$, this inequality in achieved size rate is strict.*

The proof of this lemma is presented in Appendix A and the technique, in its current form, is specific to Assumption 3. Further for Assumption 3 the following corollary, that follows from Theorem 1 and Lemma 1, states that the f-TTL achieves an average size rate which is less or equal to the target size rate.

Corollary 1. *Assume the target tuple of hit rate and size rate, $(\boldsymbol{h}^*, \boldsymbol{s}^*)$, is $(1-2\epsilon)\boldsymbol{L}$-feasible with $\epsilon > 0$ and additionally, $\boldsymbol{h}^*$ is 'typical'. Under Assumption 3 with $\|\boldsymbol{L}\|_\infty$-rarity condition, a f-TTL algorithm with parameters $\boldsymbol{L}$, ϵ, achieves asymptotically almost surely a tuple $(\boldsymbol{h}^*, \boldsymbol{s})$ with size rate $\boldsymbol{s} \preccurlyeq \boldsymbol{s}^*$.*

5. IMPLEMENTATION OF D- AND F-TTL

One of the main practical challenges in implementing d-TTL and f-TTL is adapting θ and θ^s to achieve the desired hit rate in the presence of unpredictable non-stationary traffic. We observe the following major differences between the theoretical and practical settings. First, the arrival process in practice changes over time (e.g. day-night variations) whereas our model assumes the stationary part is fixed. Second, in practice performance in finite time horizons is of particular interest. But the effect of non-stationary traffic in a finite time window is not negligible under our model which matches reality (flash crowd in specific time windows). We now discuss some modifications we make to translate theory to practice and evaluate these modification in Section 6.

Fixing the maximum TTL. The truncation parameter (maximum TTL value) $\boldsymbol{L}$ defined in Section 3 is crucial in the analysis of the system. However, in practice, we can choose an arbitrarily large value such that we let θ explore a larger space to achieve the desired hit rate in both d-TTL and f-TTL.

Constant step sizes for θ and θ^s updates. Algorithms 1 and 2 use decaying step sizes $\eta(l)$ and $\eta_s(l)$ while adapting θ and θ^s. This is not ideal in practical settings where the traffic composition is constantly changing, and we need θ and θ^s to capture those variations. Therefore, we choose carefully hand-tuned constant step sizes that capture the variability in traffic well. We discuss the sensitivity of the d-TTL and f-TTL algorithms to changes in the step size in Section 6.6. Furthermore, it is possible to extend the convergence results to this setting with 'high' probability guarantees (See [26]).

Tuning size rate targets. In practice, f-TTL may not be able to achieve small size rate targets in the presence of time varying and non-negligible non stationary traffic. In such cases, f-TTL uses the size rate target to induce filtering. For instance, when there is a sudden surge of non-stationary content, θ^s can be aggressively reduced by setting small size rate targets. This in turn filters out a lot of non-stationary objects while an appropriate increase in θ maintains the target hit rate. Hence, the size rate target can be used as a tunable knob in CDNs to adaptively filter out unpredictable non-stationary content. In our experiments in Section 6, we use a target size rate that is 50% of the size rate of d-TTL. This forces f-TTL to achieve the same hit rate as d-TTL but at half the cache space, if feasible. It should be noted that a size rate target of 0 while most aggressive, is not necessarily the best target. This is because, size rate target of 0, sets θ^s to 0 and essentially increases the θ to attain the hit rate target. This may lead to an increase in average cache size when compared to an f-TTL implementation with non-zero size rate target. Specifically, in our simulations at 40% target OHR setting non-zero size rate target leads to nearly 15% decrease in the average cache size as compared to the zero size rate target.

6. EMPIRICAL EVALUATION

We evaluate the performance of d-TTL and f-TTL implementations, both in terms of the hit rate achieved and the cache size requirements, using actual production traces from a major CDN.

6.1 Experimental setup

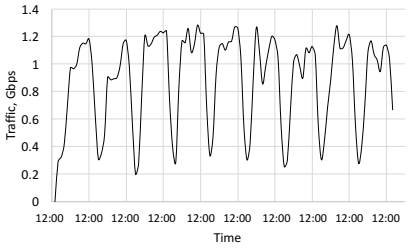

Figure 1: Content traffic served to users from the CDN server, averaged every 2 hours. The traces were collected from 29^{th} January to 6^{th} February 2015.

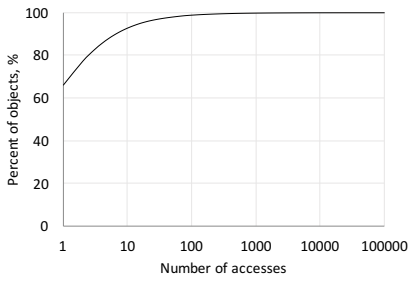

Figure 2: Popularity of content accessed by users in the 9-day period.

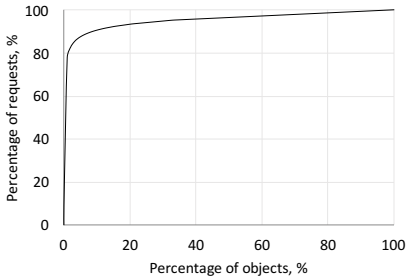

Figure 3: A large fraction of the requests are for a small fraction of the objects.

Content Request Traces. We use an extensive data set containing access logs for content requested by users that we collected from a typical production server in Akamai's commercially-deployed CDN [2]. Each log line corresponds to a single request and contains a timestamp, the requested URL (anonymized), object size, and bytes served for the request. The access logs were collected over a period of 9 days. The traffic served in Gbps captured in our data set is shown in Figure 1. We see that there is a diurnal traffic pattern with the first peak generally around 12PM (probably due to increased traffic during the afternoon) and the second peak occurring around 10-11PM. There is a small dip in traffic during the day between 4-6PM. This could be during evening commute when there is less internet traffic. The lowest traffic is observed at the early hours of the morning between 4AM and 9AM.

The content requests traces used in this work contain 504 million requests (resp., 165TB) for 25 million distinct objects (resp., 15TB). As shown in Figure 2, the content popularity distribution exhibits a "long tail" with nearly 70% of the objects accessed only once. The presence of a significant amount of non-stationary traffic in the form of "one-hit-wonders" in production traffic is consistent with similar observations made in earlier work [17]. Further, as shown in Figure 3, 80% of the requests are for 1% of the objects.

Trace-based Cache Simulator. We built a custom event-driven simulator to simulate the different TTL caching algorithms. The simulator takes as input the content traces and computes a number of statistics such as the hit rate obtained over time, the variation in θ, θ^s and the cache size over time. We implement and simulate both d-TTL and f-TTL using the parameters listed in Table 1. We use constant step sizes, η=1e-2 and η_s=1e-9, while adapting the values of θ and θ^s. The values chosen were found to capture the variability in our input trace well. We evaluate the sensitivity of d-TTL and f-TTL to changes in η and η_s in Section 6.6.

Simulation length	9 days
Minimum TTL value	0 secs
Maximum TTL value	10^7 secs
Step size for θ updates	η
Step size for θ^s updates	$\frac{\eta_s}{\text{target size rate} \times \text{average object size}}$

Table 1: Simulation parameters.

6.2 How TTLs adapt over time

To understand how d-TLL and f-TTL adapt their TTLs over time in response to request dynamics, we simulated these algorithms with a target object hit rate of 60% and a target size rate that is 50% of the size rate achieved by d-TTL. In Figure 4 we plot the traffic in Gbps, the variation in θ for d-TTL, θ for f-TTL and θ^s over time, all averaged over 2 hour windows. We consider only the object hit rate scenario to explain the dynamics. We observe similar performance when we consider byte hit rates.

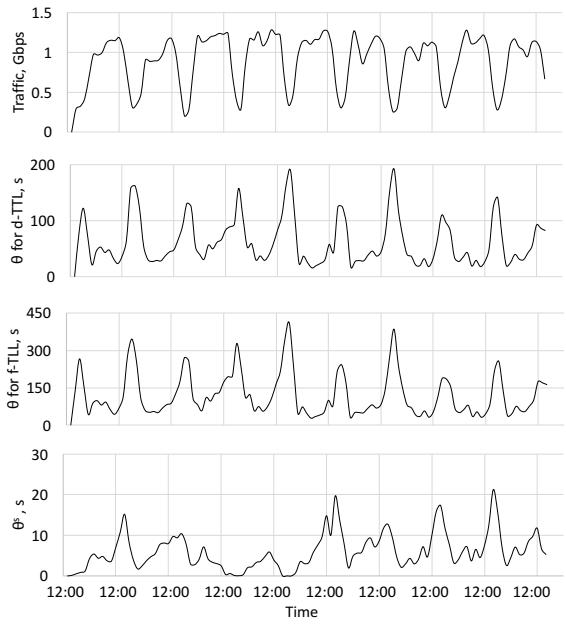

Figure 4: Variation in θ for d-TTL, θ for f-TTL and θ^s over time with target object hit rate=60%.

From Figure 4, we see that the value of θ for d-TTL is smaller than that of f-TTL. This happens due to the fact that f-TTL filters out rare objects to meet the target size rate, which can in turn reduce the hit rate, resulting in an increase in θ to achieve the target hit rate. We also observe that θ for both d- and f-TTL is generally smaller during peak hours when compared to off-peak hours. This is because the inter-arrival time of popular content is smaller during peak

hours. Hence, a smaller θ is sufficient to achieve the desired hit rate. However, during off-peak hours, traffic drops by almost 70%. With fewer content arrivals per second, θ increases to provide the same hit rate. In the case of f-TTL, the increase in θ, increases the size rate of the system, which in turn leads to a decrease in θ^s. This matches with the theoretical intuition that d-TTL adapts θ only to achieve the target hit rate while f-TTL adapts both θ and θ^s to reduce the cache size while also achieving the target hit rate.

6.3 Hit rate performance of d-TTL and f-TTL

The performance of a caching algorithm is often measured by its hit rate curve (HRC) that relates its cache size with the (object or byte) hit rate that it achieves. HRCs are useful for CDNs as they help provision the right amount of cache space to obtain a certain hit rate. We compare the HRCs of d-TTL and f-TTL for object hit rates and show that f-TTL significantly outperforms d-TTL by filtering out the rarely-accessed non-stationary objects. The HRCs for byte hit rates are shown in Appendix B.

To obtain the HRC for d-TTL, we fix the target hit rate at 80%, 70%, 60%, 50% and 40% and measure the hit rate and cache size achieved by the algorithm. Similarly, for f-TTL, we fix the target hit rates at 80%, 70%, 60%, 50% and 40%. Further, we set the target size rate f-TTL to 50% of the size rate of d-TTL. The HRCs for object hit rates are shown in Figures 5. The hit rate performance for byte hit rates is discussed in Appendix B. Note that the y-axis is presented in log scale for clarity.

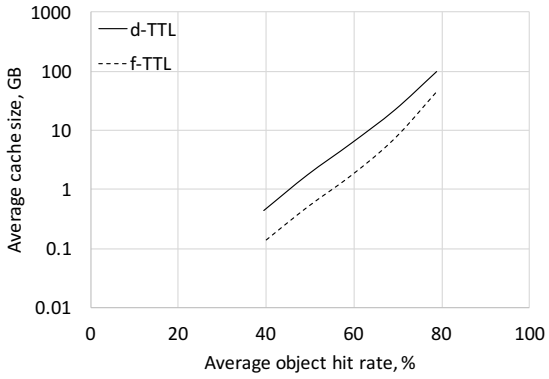

Figure 5: Hit rate curve for object hit rates.

From Figure 5 we see that f-TTL always performs better than d-TTL i.e. for a given hit rate, f-TTL requires lesser cache size on average than d-TTL. In particular, on average, f-TTL requires a cache that is 49% smaller than d-TTL to achieve the same object hit rate.

6.4 Convergence of d-TTL and f-TTL

For the dynamic TTL algorithms to be useful in practice, they need to converge to the target hit rate with low error. In this section we measure the object hit rate convergence over time, averaged over the entire time window and averaged over 2 hour windows for both d-TTL and f-TTL. We set the target hit rate to 60% and a target size rate that is 50% of the size rate of d-TTL. The byte hit rate convergence is discussed in Appendix B.

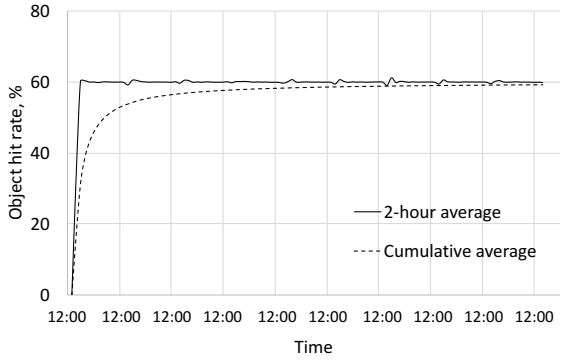

Figure 6: Object hit rate convergence over time for d-TTL; target object hit rate=60%.

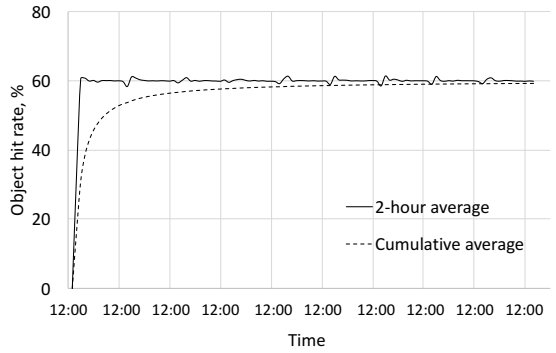

Figure 7: Object hit rate convergence over time for f-TTL; target object hit rate=60%.

From Figures 6 and 7, we see that 2 hour averaged object hit rates achieved by both d-TTL and f-TTL have a cumulative error of less than 1.3% while achieving the target object hit rate on average. We see that both d-TTL and f-TTL tend to converge to the target hit rate, which illustrates that both d-TTL and f-TTL are able to adapt well to the dynamics of the input traffic.

In general, we also see that d-TTL has lower variability for object hit rate compared to f-TTL due to the fact that d-TTL does not have any bound on the size rate while achieving the target hit rate, while f-TTL is constantly filtering out non-stationary objects to meet the target size rate while also achieving the target hit rate.

6.5 Accuracy of d-TTL and f-TTL

A key goal of the dynamic TTL algorithms (d-TTL and f-TTL) is to achieve a target hit rate, even in the presence of bursty and non-stationary requests. We evaluate the performance of both these algorithms by fixing the target hit rate and comparing the hit rates achieved by d-TTL and f-TTL with caching algorithms such as Fixed TTL (TTL-based caching algorithm that uses a constant TTL value) and LRU (constant cache size), provisioned using Che's approximation [10]. We only present the results for object hit rates (OHR), similar behavior is observed for byte hit rates. The results are shown in Table 2.

For this evaluation, we fix the target hit rates (column 1)

Target	Fixed TTL (Che's approx.)			LRU (Che's approx.)		d-TTL		f-TTL	
OHR (%)	TTL (s)	OHR (%)	Size (GB)	OHR (%)	Size (GB)	OHR (%)	Size (GB)	OHR (%)	Size (GB)
80	2784	83.29	217.11	84.65	316.81	78.72	97.67	78.55	55.08
70	554	75.81	51.88	78.37	77.78	69.21	21.89	69.14	11.07
60	161	68.23	16.79	71.64	25.79	59.36	6.00	59.36	2.96
50	51	60.23	5.82	64.18	9.2	49.46	1.76	49.47	0.86
40	12	50.28	1.68	54.29	2.68	39.56	0.44	39.66	0.20

Table 2: Comparison of target hit rate and average cache size achieved by d-TTL and f-TTL with Che's approximation, Fixed-TTL and LRU.

Target	Average OHR (%)			Average cache size (GB)			5% outage fraction		
OHR (%)	$\eta = 0.1$	$\eta = 0.01$	$\eta = 0.001$	$\eta = 0.1$	$\eta = 0.01$	$\eta = 0.001$	$\eta = 0.1$	$\eta = 0.01$	$\eta = 0.001$
60	59.35	59.36	59.17	9.03	6.00	5.41	0.01	0.01	0.05
80	79.13	78.72	77.69	150.56	97.67	75.27	0.07	0.11	0.23

Table 3: Impact of exponential changes in constant step size η on the average hit rate, average cache size and 5% outage fraction of d-TTL (sensitivity analysis of d-TTL).

Target	Average OHR (%)			Average cache size (GB)			5% outage fraction		
OHR (%)	η(1+0.05)	η	η(1-0.05)	η(1+0.05)	η	η(1-0.05)	η(1+0.05)	η	η(1-0.05)
60	59.36	59.36	59.36	5.98	6.00	6.02	0.01	0.01	0.01
80	78.73	78.72	78.71	98.21	97.67	97.1	0.11	0.11	0.11

Table 4: Impact of linear changes in constant step size $\eta = 0.01$ on the average hit rate, average cache size and 5% outage fraction of d-TTL (robustness analysis of d-TTL).

Target	Average OHR (%)			Average cache size (GB)			5% outage fraction		
OHR (%)	η_s = 1e-8	η_s = 1e-9	η_s = 1e-10	η_s = 1e-8	η_s = 1e-9	η_s = 1e-10	η_s = 1e-8	η_s = 1e-9	η_s = 1e-10
60	59.36	59.36	59.36	5.46	2.96	1.88	0.01	0.01	0.02
80	78.65	78.55	78.47	89.52	55.08	43.34	0.12	0.14	0.17

Table 5: Impact of exponential changes in constant step size η_s on the average hit rate, average cache size and 5% outage fraction of f-TTL (sensitivity analysis of f-TTL).

Target	Average OHR (%)			Average cache size (GB)			5% outage fraction		
OHR (%)	η_s(1+0.05)	η_s	η_s(1-0.05)	η_s(1+0.05)	η_s	η_s(1-0.05)	η_s(1+0.05)	η_s	η_s(1-0.05)
60	59.36	59.36	59.36	3.01	2.96	2.91	0.01	0.01	0.01
80	78.55	78.55	78.54	55.65	55.08	54.27	0.14	0.14	0.14

Table 6: Impact of linear changes in constant step size η_s = 1e-9 on the average hit rate, average cache size and 5% outage fraction of f-TTL (robustness analysis of f-TTL).

and compute the TTL and cache size using Che's approximation (columns 2 and 6). We then measure the hit rate and cache size of Fixed TTL (columns 3 and 4) using the TTL computed in column 2, and the hit rate of LRU (column 5) using the cache size computed in column 6. Finally, we compute the hit rate and cache size achieved by d-TTL and f-TTL (columns 7-10) to achieve the target hit rates in column 1 and a target size rate that is 50% of that of d-TTL.

We make the following conclusions from Table 2.

1. The d-TTL and f-TTL algorithms meet the target hit rates with a small error of 1.2% on average. This is in contrast to the Fixed TTL algorithm which has a high error of 14.4% on average and LRU which has even higher error of 20.2% on average. This shows that existing algorithms such as Fixed TTL and LRU are unable to meet the target hit rates while using heuristics such as Che's approximation, which cannot account for non-stationary content.

2. The cache size required by d-TTL and f-TTL is 23.5% and 12% respectively, of the cache size estimated by Che's approximation and 35.8% and 18.3% respectively, of the cache size achieved by the Fixed TTL algorithm, on average. This indicates that both LRU and the Fixed TTL algorithm, provisioned using Che's approximation, grossly overestimate the cache size requirements.

6.6 Sensitivity and robustness of d-TTL and f-TTL

We use constant step sizes while adapting the values of θ and θ_s in practical settings for reasons discussed in Section 5. In this section, we evaluate the sensitivity and robustness of d-TTL and f-TTL to the chosen step sizes. The sensitivity captures the change in performance due to large changes in step size, whereas the robustness captures the change due to small perturbations around a specific step size. For ease of explanation, we only focus on two target object hit rate, 60% and 80% corresponding to medium and high hit rates. The observations are similar for other target hit rates and for byte hit rates.

Table 3 illustrates the sensitivity of d-TTL to exponential

changes in the step size η. For each target hit rate, we measure the average hit rate achieved by d-TTL, the average cache size and the 5% outage fraction, for each value of step size. The 5% outage fraction is defined as the fraction of time the hit rate achieved by d-TTL differs from the target hit rate by more than 5%. From this table, we see that a step size of 0.01 offers the best trade-off among the three parameters, namely average hit rate, average cache size and 5% outage fraction. Table 4 illustrates the robustness of d-TTL to small changes in the step size. We evaluate d-TTL at step sizes $\eta = 0.01 \times (1 \pm 0.05)$. We see that d-TTL is robust to small changes in the step size.

To evaluate the sensitivity and robustness of f-TTL, we fix the step size $\eta = 0.01$ to update θ and evaluate the performance of f-TTL at different step sizes, η_s, to update θ_s. The results for sensitivity and robustness are shown in Tables 5 and 6 respectively. For f-TTL, we see that a step size of η_s=1e-9 offers the best tradeoff among the different parameters namely average hit rate, average cache size and 5% outage fraction. Like, d-TTL, f-TTL is robust to the changes in step size parameter η_s.

7. RELATED WORK

Caching algorithms have been studied for decades in different contexts such as CPU caches, memory caches, CDN caches and so on. We briefly review some prior work that is relevant to this paper.

TTL-based caching. TTL caches have found their place in theory as a tool for analyzing capacity based caches [9, 8, 5], starting from Che's Approximation of LRU caches [10]. Recently, its generalizations [13, 27] have commented on its wider applicability. However, the generalizations hint towards the need for more tailored approximations and show that the vanilla Che's approximation can be inaccurate [27]. On the applications side, recent works have demonstrated the use of TTL caches in utility maximization [9] and hit ratio maximization [7]. Specifically, in [9] the authors provide an online TTL adaptation highlighting the need for adaptive algorithms. However, unlike prior work, we propose the first adaptive TTL-based caching algorithms that provides provable hit rate and size rate performance in the presence of non-stationary traffic such as one-hit wonders and traffic bursts.

We also review some capacity-based caching algorithms.

Capacity-based caching. Capacity-based caching algorithms have been in existence for over 4 decades and have been studied both theoretically (e.g. exact expressions [28, 29] and mean field approximations [30, 31, 32] for hit rates) and empirically (in the context of web caching: [33]). Various cache replacement algorithms have been proposed based on the frequency of object requests (e.g. LFU), recency of object requests (e.g. LRU) or a combination of the two parameters (e.g., LRU-K, 2Q, LRFU [34, 23, 35]). Given that CDNs see a lot of non-stationary traffic, cache admission policies such as those using bloom filters [19] have also been proposed to maximize the hit rate under space constraints. Further, non-uniform distribution of object sizes have led to more work that admit objects based on the size (e.g., LRU-S, AdaptSize [36, 37]). While several capacity-based algorithms have been proposed, most don't provide theoretical performance guarantees in achieving target hit rates.

Cache Tuning and Adaptation. Most existing adaptive caching algorithms require careful parameter tuning to work in practice. There have been two main cache tuning methods: (1) global search over parameters based on prediction model, e.g. [38], and (2) simulation and parameter optimization based on shadow cache, e.g. [39]. The first method often fails in the presence of cache admission policyes; whereas, the second method typically assumes stationary arrival processes to work well. However, with real traffic, static parameters are not desirable [22] and an adaptive/self-tuning cache is necessary. The self-tuning heuristics include, e.g., ARC [22], CAR [40], PB-LRU [41], which try to adapt cache partitions based on system dynamics. While these tuning methods are meant to deal with non-stationary traffic, they lack theoretical guarantees unlike our work, where we provably achieve a target hit rate and a feasible size rate by dynamically changing the TTLs of cached content.

Finally, we discuss work related to cache hierarchies, highlighting the differences between those and the proposed f-TTL algorithm.

Cache hierarchies. Cache hierarchies, made popular for web caches in [42, 10, 43], consist of separate caches, mostly LRU [44, 43] or TTL-based [5], arranged in multiple levels; with users at the lowest level and the server at the highest. A requested object is fetched from the lowest possible cache and, typically, replicated in all the caches on the request path. Initial work [5] present analysis for lines and trees of caches, which is later generalized to more complex networks [44] and in the presence of renewal arrivals [6].

While similar in spirit, the f-TTL algorithm differs in its structure and operation from hierarchical caches. Besides the adaptive nature of the TTLs, the higher and lower-level caches are assumed to be co-located and no object is replicated between them—a major structural and operational difference. Further, the lower-level cache uses a smaller TTL to filter out rare and unpopular content, similar to caching algorithms such as 2Q [23], bloom filters [19], to name a few. Also, the lower-level cache maintains a shadow cache to *remember* accesses to rare content, thereby adapting the TTL to achieve the target hit rate.

8. CONCLUSIONS

In this paper we designed adaptive TTL caching algorithms that can automatically learn and adapt to the request traffic and provably achieve any feasible hit rate and cache size. Our work fulfills a long-standing deficiency in the modeling and analysis of caching algorithms in presence of bursty and non-stationary request traffic. In particular, we presented a theoretical justification for the use of two-level caches in CDN settings where large amounts of non-stationary traffic can be filtered out to conserve cache space while also achieving target hit rates. On the practical side, we evaluated our TTL caching algorithms using traffic traces from a production Akamai CDN server. The evaluation results show that our adaptive TTL algorithms can achieve the target hit rate with high accuracy; further, the two-level TTL algorithm can achieve the same target hit rate at much smaller cache size.

9. REFERENCES

[1] John Dilley, Bruce Maggs, Jay Parikh, Harald Prokop, Ramesh Sitaraman, and Bill Weihl. Globally distributed content delivery. *Internet Computing, IEEE*, 6(5):50–58, 2002. http://www.computer.org/internet/ic2002/w5050abs.htm.

[2] E. Nygren, R.K. Sitaraman, and J. Sun. The Akamai Network: A platform for high-performance Internet applications. *ACM SIGOPS Operating Systems Review*, 44(3):2–19, 2010.

[3] Cisco. Visual Networking Index: The Zettabyte Era—Trends and Analysis, june 2016. goo.gl/id2RB4.

[4] Jaeyeon Jung, Arthur W Berger, and Hari Balakrishnan. Modeling ttl-based internet caches. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 1, pages 417–426. IEEE, 2003.

[5] N Choungmo Fofack, Philippe Nain, Giovanni Neglia, and Don Towsley. Analysis of ttl-based cache networks. In *Performance Evaluation Methodologies and Tools (VALUETOOLS), 2012 6th International Conference on*, pages 1–10. IEEE, 2012.

[6] Daniel S Berger, Philipp Gland, Sahil Singla, and Florin Ciucu. Exact analysis of ttl cache networks. *Performance Evaluation*, 79:2–23, 2014.

[7] Daniel S Berger, Sebastian Henningsen, Florin Ciucu, and Jens B Schmitt. Maximizing cache hit ratios by variance reduction. *ACM SIGMETRICS Performance Evaluation Review*, 43(2):57–59, 2015.

[8] Michele Garetto, Emilio Leonardi, and Valentina Martina. A unified approach to the performance analysis of caching systems. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, 1(3):12, 2016.

[9] Mostafa Dehghan, Laurent Massoulie, Don Towsley, Daniel Menasche, and YC Tay. A utility optimization approach to network cache design. *arXiv preprint arXiv:1601.06838*, 2016.

[10] Hao Che, Ye Tung, and Zhijun Wang. Hierarchical web caching systems: Modeling, design and experimental results. *IEEE Journal on Selected Areas in Communications*, 20(7):1305–1314, 2002.

[11] Christine Fricker, Philippe Robert, James Roberts, and Nada Sbihi. Impact of traffic mix on caching performance in a content-centric network. In *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*, pages 310–315. IEEE, 2012.

[12] Christine Fricker, Philippe Robert, and James Roberts. A versatile and accurate approximation for lru cache performance. In *Proceedings of the 24th International Teletraffic Congress*, page 8. International Teletraffic Congress, 2012.

[13] Giuseppe Bianchi, Andrea Detti, Alberto Caponi, and Nicola Blefari Melazzi. Check before storing: what is the performance price of content integrity verification in lru caching? *ACM SIGCOMM Computer Communication Review*, 43(3):59–67, 2013.

[14] Fabrice Guillemin, Bruno Kauffmann, Stephanie Moteau, and Alain Simonian. Experimental analysis of caching efficiency for youtube traffic in an isp network. In *Teletraffic Congress (ITC), 2013 25th International*, pages 1–9. IEEE, 2013.

[15] Emilio Leonardi and Giovanni Luca Torrisi. Least recently used caches under the shot noise model. In *Computer Communications (INFOCOM), 2015 IEEE Conference on*, pages 2281–2289. IEEE, 2015.

[16] Nicolas Gast and Benny Van Houdt. Asymptotically exact ttl-approximations of the cache replacement algorithms lru (m) and h-lru. In *Teletraffic Congress (ITC 28), 2016 28th International*, volume 1, pages 157–165. IEEE, 2016.

[17] Bruce M. Maggs and Ramesh K. Sitaraman. Algorithmic nuggets in content delivery. *SIGCOMM Comput. Commun. Rev.*, 45(3):52–66, July 2015.

[18] Vijay R Konda and John N Tsitsiklis. Onactor-critic algorithms. *SIAM journal on Control and Optimization*, 42(4):1143–1166, 2003.

[19] Bruce M Maggs and Ramesh K Sitaraman. Algorithmic nuggets in content delivery. *ACM SIGCOMM Computer Communication Review*, 45(3):52–66, 2015.

[20] Jaeyeon Jung, Balachander Krishnamurthy, and Michael Rabinovich. Flash crowds and denial of service attacks: Characterization and implications for cdns and web sites. In *Proceedings of the 11th international conference on World Wide Web*, pages 293–304. ACM, 2002.

[21] Bennett Fox. Semi-markov processes: A primer. Technical report, DTIC Document, 1968.

[22] Nimrod Megiddo and Dharmendra S Modha. Outperforming lru with an adaptive replacement cache algorithm. *Computer*, 37(4):58–65, 2004.

[23] Theodore Johnson and Dennis Shasha. X3: A low overhead high performance buffer management replacement algorithm. 1994.

[24] Harold Kushner and G George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.

[25] Shalabh Bhatnagar, Richard S Sutton, Mohammad Ghavamzadeh, and Mark Lee. Natural actor–critic algorithms. *Automatica*, 45(11):2471–2482, 2009.

[26] Vladislav B Tadic and Sean P Meyn. Asymptotic properties of two time-scale stochastic approximation algorithms with constant step sizes. In *American Control Conference, 2003. Proceedings of the 2003*, volume 5, pages 4426–4431. IEEE, 2003.

[27] Felipe Olmos, Bruno Kauffmann, Alain Simonian, and Yannick Carlinet. Catalog dynamics: Impact of content publishing and perishing on the performance of a lru cache. In *Teletraffic Congress (ITC), 2014 26th International*, pages 1–9. IEEE, 2014.

[28] Erol Gelenbe. A unified approach to the evaluation of a class of replacement algorithms. *IEEE Transactions on Computers*, 100(6):611–618, 1973.

[29] Oleg Ivanovich Aven, Edward Grady Coffman, and Yakov Afroimovich Kogan. *Stochastic analysis of computer storage*, volume 38. Springer Science & Business Media, 1987.

[30] Ryo Hirade and Takayuki Osogami. Analysis of page replacement policies in the fluid limit. *Operations research*, 58(4-part-1):971–984, 2010.

[31] Naoki Tsukada, Ryo Hirade, and Naoto Miyoshi. Fluid limit analysis of fifo and rr caching for independent reference models. *Performance Evaluation*, 69(9):403–412, 2012.

[32] Nicolas Gast and Benny Van Houdt. Transient and

steady-state regime of a family of list-based cache replacement algorithms. *ACM SIGMETRICS Performance Evaluation Review*, 43(1):123–136, 2015.

[33] Stefan Podlipnig and Laszlo Böszörmenyi. A survey of web cache replacement strategies. *ACM Computing Surveys (CSUR)*, 35(4):374–398, 2003.

[34] Elizabeth J O'neil, Patrick E O'neil, and Gerhard Weikum. The lru-k page replacement algorithm for database disk buffering. *ACM SIGMOD Record*, 22(2):297–306, 1993.

[35] Donghee Lee, Jongmoo Choi, Jong-Hun Kim, Sam H Noh, Sang Lyul Min, Yookun Cho, and Chong Sang Kim. On the existence of a spectrum of policies that subsumes the least recently used (lru) and least frequently used (lfu) policies. In *ACM SIGMETRICS Performance Evaluation Review*, volume 27, pages 134–143. ACM, 1999.

[36] David Starobinski and David Tse. Probabilistic methods for web caching. *Performance evaluation*, 46(2):125–137, 2001.

[37] Daniel S Berger, Ramesh K Sitaraman, and Mor Harchol-Balter. Achieving high cache hit ratios for cdn memory caches with size-aware admission. 2016.

[38] Asaf Cidon, Assaf Eisenman, Mohammad Alizadeh, and Sachin Katti. Dynacache: Dynamic cloud caching. In *7th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 15)*, 2015.

[39] Asaf Cidon, Assaf Eisenman, Mohammad Alizadeh, and Sachin Katti. Cliffhanger: scaling performance cliffs in web memory caches. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pages 379–392, 2016.

[40] Sorav Bansal and Dharmendra S Modha. Car: Clock with adaptive replacement. In *FAST*, volume 4, pages 187–200, 2004.

[41] Qingbo Zhu, Asim Shankar, and Yuanyuan Zhou. Pb-lru: a self-tuning power aware storage cache replacement algorithm for conserving disk energy. In *Proceedings of the 18th annual international conference on Supercomputing*, pages 79–88. ACM, 2004.

[42] Anawat Chankhunthod, Peter B Danzig, Chuck Neerdaels, Michael F Schwartz, and Kurt J Worrell. A hierarchical internet object cache. In *USENIX Annual Technical Conference*, pages 153–164, 1996.

[43] Ramesh K. Sitaraman, Mangesh Kasbekar, Woody Lichtenstein, and Manish Jain. Overlay networks: An Akamai perspective. In *Advanced Content Delivery, Streaming, and Cloud Services*. John Wiley & Sons, 2014.

[44] Elisha J Rosensweig, Jim Kurose, and Don Towsley. Approximate models for general cache networks. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9. IEEE, 2010.

[45] Walter L Smith. Renewal theory and its ramifications. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 243–302, 1958.

[46] Sean P Meyn and Richard L Tweedie. *Markov chains and stochastic stability*. Springer Science & Business Media, 2012.

[47] Vladislav B Tadic. Almost sure convergence of two time-scale stochastic approximation algorithms. In *American Control Conference, 2004. Proceedings of the 2004*, volume 4, pages 3802–3807. IEEE, 2004.

[48] Irwin Sandberg. Global implicit function theorems. *IEEE Transactions on Circuits and Systems*, 28(2):145–149, 1981.

APPENDIX

A. PROOF OF MAIN RESULTS

The proof of the main theorem is separated in three main parts. In the first part we formally introduce the dynamic system as a discrete time Markov process on continuous state space, a.k.a. Harris chains. Then we shift our focus towards analyzing the static caching system when the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^s$ are assumed to be fixed. In the last part, we state and prove the sufficient conditions for the convergence of the two timescale stochastic approximation algorithm. Then in the final subsection we provide proofs for the special cases of Poisson arrivals with independent labeling and sub-classes of that.

In the proof a scalar function when presented in boldface means we apply the same function to each coordinate of its vector argument, i.e. $\boldsymbol{f}(\boldsymbol{x}) = (f(x_1), \cdots, f(x_n))$. Further, $\mathbf{1}_n$ denotes the all 1 vector of size n and $\mathbf{e}_i^n$ denotes the i-th standard basis for $\mathbb{R}^n$ (i-th coordinate is equal to 1 and other coordinates are equal to 0).

A.1 Caching Process

In this section we study the evolution of the caching system as a discrete time Markov chain on a continuous state space, commonly known as Harris chains. We first observe that, by fixing $\boldsymbol{\vartheta}^s = \mathbf{1}$ the f-TTL quickly identifies with the d-TTL system (when we treat $\boldsymbol{\theta}^s = \boldsymbol{\theta}$). We call it f-TTL cache collapsing to d-TTL cache.[6] This lead us to unify the two systems, d-TTL and f-TTL system, where for d-TTL $\boldsymbol{\vartheta}^s(l) = \mathbf{1}$ for all l.

We next develop the necessary notations for the rest of the proofs before formally describing the system evolution.

Arrival Process Notation. Let the sequence of arrivals in the system be given by $\boldsymbol{A} = \{A(l) : l \in \mathbb{N}\}$ where each inter arrival time $X(l) = A(l) - A(l-1)$ has identical distribution. By convention define $A(0) = 0$. Prior works on TTL algorithm with a fixed TTL consider i.i.d. inter arrival times for each object and give the average cache hit rate, i.e. the probability a cache hit occurs for an incoming request. In this paper we consider a more general request model where the request labels have Markovian dependence across subsequent arrivals. Let $L_c = \{l : c(l) = c, l \in \mathbb{N}\}$ denote the sequence of arrivals where object c is requested, for $c \in \mathcal{U}$.

The request arrival times of a object c is given as $\boldsymbol{A}_c = \{A(l) : l \in L_c\}$ and $A_c(j)$ denotes the j-th arrival of object c. Denote the j-th inter arrival time for object c as $X_c(j) = A_c(j) - A_c(j-1)$. For notational convenience let the inter arrival distribution for object c be $p_c(\cdot) = \mathbb{P}(X_c(1) \leq \cdot)$ and its mean be $1/\lambda_c$. Further by $N_c(u) = |\{j : A_c(j) \leq u\}|$

[6]Technically, the f-TTL with $\boldsymbol{\vartheta}^s = \mathbf{1}$ gets coupled to d-TTL from time $\tau(l) = \|\boldsymbol{L}\|_\infty = L_{max}$ onwards with the d-TTL system. Further, any sample path (with non zero probability) for the d-TTL system is contained in at least one sample path (with non zero probability) of collapsed f-TTL. Therefore, any a.s. convergence in the collapsed f-TTL will imply an a.s. convergence in the d-TTL.

denote the number of arrival of object c up to time u. Also call the realization of the process Z as $\{Z(l) : l \in \mathbb{N}\}$ and let $Z(0)$ be the initial state. Specifically, $Z(l)$ denotes the $(l-1)$-th request.

System State Notation. The evolution of the system is a coupled stochastic process given as the sequence, $\{(\boldsymbol{\vartheta}(l), \boldsymbol{\vartheta}^s(l), \mathcal{S}(l) \equiv (\boldsymbol{\Psi}(l), Z(l))) : l \in \mathbb{N}\}$. At the l-th arrival, $\boldsymbol{\Psi}(l)$ represents the TTL information of the recurrent objects and $Z(l)$ denotes which object (object-type for 'rare' objects) is requested in the $(l-1)$-th arrival. The vector of timer tuples for *recurring objects* is $\boldsymbol{\Psi}(l) = \left\{\psi_c(l) \equiv (\psi_c^0(l), \psi_c^1(l), \psi_c^2(l)) : c \in \mathcal{K}\right\}$. We intentionally loose track of the *rare objects* except for the information about $Z(l)$ while maintaining system state. Finally, let $\mathcal{F}_l$ be a sequence of non decreasing σ-algebras such that the history of the system adaptation upto the l-th arrival, $\mathcal{H}(l) = \{\theta(0), Z(0), \mathcal{S}(i), X(i-1) : i \le l\}$ is $\mathcal{F}_l$-measurable.

Evolution. We concisely represent the evolution of $\boldsymbol{\vartheta}(l)$ and $\boldsymbol{\theta}(l)$ as (3),

$$\begin{aligned}\boldsymbol{\vartheta}(l+1) &= \mathcal{P}_H\left(\boldsymbol{\vartheta}(l) + \eta(l)\widehat{w}(l)\left(h^*_{t(l)} - Y(l)\right)\mathbf{e}^T_{t(l)}\right).\\ \boldsymbol{\theta}(l+1) &= \sum_t L_t\vartheta_t(l+1)\mathbf{e}_t^T.\end{aligned} \tag{3}$$

Here $\eta(l) = \frac{\eta_0}{l^\alpha}$ is a decaying step size for $\alpha \in (1/2, 1)$, $\mathcal{P}_H$ denotes the projection over the set $H = \{\mathbf{x} \in \mathbb{R}^T : 0 \le x_i \le 1\}$, $\boldsymbol{L}$ is the upper bound of the TTL which is provided as an input, and

$$\widehat{w}(l) = \begin{cases} w(l), & \text{for byte hit rate,}\\ 1, & \text{for object hit rate.}\end{cases}$$

Further, the evolution of $\boldsymbol{\vartheta}^s(l)$ and $\boldsymbol{\theta}^s(l)$ is represented as (4).

$$\begin{aligned}\boldsymbol{\vartheta}^s(l+1) &= \mathcal{P}_H\left(\boldsymbol{\vartheta}^s(l) + \eta_s(l)w(l)(s^*_{t(l)} - s(l))\mathbf{e}^T_{t(l)}\right)\\ \boldsymbol{\theta}^s(l+1) &= \sum_t L_t\vartheta_t(l+1)\Gamma\left(\vartheta_t(l+1), \vartheta^s_t(l+1); \epsilon\right)\mathbf{e}_t^T.\end{aligned} \tag{4}$$

Here $\eta_s(l) = \eta_0/l$ and ϵ is a parameter of the algorithm. The projection $\mathcal{P}_H$ and $\boldsymbol{L}$ are defined as in d-TTL. The $\Gamma(\cdot,\cdot;\epsilon)$ is a *threshold function*, defined earlier as, a twice differentiable non-decreasing function, $\Gamma(x, y; \epsilon) : [0,1]^2 \to [0,1]$, which takes value 1 for $x \ge 1 - \epsilon/2$ and takes value y for $x \le 1 - 3\epsilon/2$, and its slope in between is bounded by $4/\epsilon$.

On l-th arrival, if a cache hit or cache virtual hit occurs for object c we have $\psi_c(l+1) = (\theta_{c_t}(l), 0, 0)$ if $c = Z(l+1)$ and $c \in \mathcal{K}$. Instead if a cache miss happens for object c we have $\psi_c(l+1) = (0, \theta^s_{c_t}(l), \theta_{c_t}(l))$. The dynamics of $Z(l)$ is Markovian and the transition probability from $Z(l)$ to a object c is given as $P(Z(l), c)$, for all objects c. For notational convenience we define the cache miss indicator, $q_c^m(l)$, cache virtual hit indicator, $q_c^v(l)$ and cache hit indicator for two caches $q_c^h(l)$, $q_c^{hs}(l)$ for all $c \in \mathcal{K}$ and for all $l \in \mathbb{N}$ as given in the following equations. The cache hit event from c is given by $(q_c^h(l) + q_c^{hs}(l))$.

$$\begin{aligned}q_c^m(l) &= \mathbb{1}\left(\max\{\psi_c^0(l), \psi_c^1(l), \psi_c^2(l)\} < X(l)\right),\\ q_c^v(l) &= \mathbb{1}\left(\psi_c^1(l) < X(l) \wedge \psi_c^2(l) \ge X(l)\right),\\ q_c^h(l) = \mathbb{1}\left(\psi_c^0(l) \ge X(l)\right),\ & q_c^{hs}(l) = \mathbb{1}\left(\psi_c^1(l) \ge X(l)\right).\end{aligned}$$

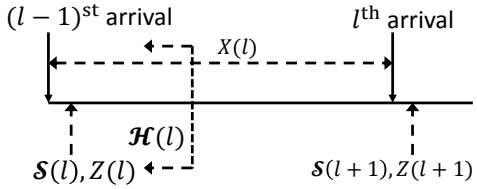

Figure 8: System Evolution

Also define the following terms for ease of presentation,

$$\begin{aligned}\tilde{p}(c,l) &= P(Z(l),c)(q_c^v(l) + q_c^h(l) + q_c^{hs}(l)),\\ \tilde{p}_1(c,l) &= P(Z(l),c)(q_c^v(l) + q_c^{hs}(l)),\\ \tilde{p}_2(c,l) &= P(Z(l),c)q_c^m(l).\end{aligned}$$

Finally, the adaptation of $\boldsymbol{\Psi}(l)$ is given as in equations (5) and (6), where $X(l)$ are i.i.d. copies with inter arrival time distribution,

$$\psi_c^0(l+1) = \begin{cases}(\psi_c^0(l) - X(l))^+ & \text{w.p. } (1 - \tilde{p}(c,l))\\ \theta_{c_t}(l) & \text{w.p. } \tilde{p}(c,l).\end{cases} \tag{5}$$

$$(\psi_c^1(l+1), \psi_c^2(l+1)) = \begin{cases}\left((\psi_c^i(l) - X(l))^+\right)_{i=1}^2 & \text{w.p. } (1 - \tilde{p}_1(c,l) - \tilde{p}_2(c,l))\\ (0,0) & \text{w.p. } \tilde{p}_1(c,l)\\ (\theta^s_{c_t}(l), \theta_{c_t}(l)) & \text{w.p. } \tilde{p}_2(c,l).\end{cases} \tag{6}$$

By $\mathbb{E}_l$, denote the conditional expectation conditioned on $\mathcal{F}(l)$, i.e. $\mathbb{E}(\cdot|\mathcal{F}(l))$. As $X(l) \sim X$, it is convenient to express the average of $q_c^{\{m,v,h,hs\}}(l)$ conditional to $\mathcal{F}_l$ as,

$$\begin{aligned}\mathbb{E}_l q_c^m(l) &= \mathbb{P}\left(X > \max\{\psi_c(l)\}\right),\\ \mathbb{E}_l q_c^v(l) &= \mathbb{P}\left(X \in (\psi_c^1(l), \psi_c^2(l)]\right),\\ \mathbb{E}_l q_c^h(l) = \mathbb{P}\left(X \le \psi_c^0(l)\right),\ & \mathbb{E}_l q_c^{hs}(l) = \mathbb{P}\left(X \le \psi_c^1(l)\right).\end{aligned}$$

$\boldsymbol{\vartheta}$ Adaptation. Recall, $Y(l)$ takes value 1 if it is a cache hit at the l-th arrival or takes the value 0 otherwise. We split $Y(l)$ in two components, $Y(l) = Y_{\mathcal{K}}(l) + \beta_h(l)$. The term $Y_{\mathcal{K}}(l) = Y(l)\mathbb{1}(c(l) \in \mathcal{K})$ denotes the contribution from *recurring objects* and $\beta_h(l) = Y(l)\mathbb{1}(c(l) \notin \mathcal{K})$ denotes the contribution from the *rare objects*. The expected (w.r.t. $\mathcal{F}(l)$) cache hit from object $c \in \mathcal{K}$ is expressed by $g_c(\cdot)$ and it is given by the following measurable function,

$$g_c(\psi_c(l), Z(l)) = \left(\mathbb{E}_l q_c^h(l) + \mathbb{E}_l q_c^{hs}(l)\right),\ \forall c \in \mathcal{K}. \tag{7}$$

Given the history $\mathcal{H}(l)$, we obtain the conditional expectation of the update of $\boldsymbol{\vartheta}(l+1)$ as (complete expression later),

$$\mathbf{g}\left(\boldsymbol{\vartheta}(l), \boldsymbol{\vartheta}^s(l), \mathcal{S}(l)\right) = \mathbb{E}_l\left(\widehat{w}(l)\left(\mathbf{h}^*_{t(l)} - Y_{\mathcal{K}}(l)\right)\mathbf{e}^T_{t(l)}\right). \tag{8}$$

Further, for all l, define the martingale difference

$$\boldsymbol{\delta M}(l) = \widehat{w(l)}\left(\mathbf{h}^*_{t(l)} - Y_{\mathcal{K}}(l)\right)\mathbf{e}^T_{t(l)} - \mathbf{g}\left(\boldsymbol{\vartheta}(l), \boldsymbol{\vartheta}^s(l), \mathcal{S}(l)\right).$$

Finally, in order to remove the projection operator we define the reflection due to this projection operator (i.e. $(\boldsymbol{x} - \mathcal{P}_H(\boldsymbol{x}))$) as $\boldsymbol{r}(l)$. We can rewrite the equation representing

the adaptation of $\boldsymbol{\theta}(l)$ in Equation (3) as;

$$\begin{aligned}\boldsymbol{\theta}(l+1) &= \mathcal{P}_H\Big(\boldsymbol{\theta}(l)+\eta(l)\mathbf{g}(\boldsymbol{\vartheta}(l),\boldsymbol{\vartheta}^s(l),\mathcal{S}(l)) \\ &\qquad +\eta(l)\left(\boldsymbol{\delta M}(l)-\beta_h(l)\mathbf{e}^T_{t(l)}\right)\Big) \\ &= \boldsymbol{\theta}(l)+\eta(l)\mathbf{g}(\boldsymbol{\vartheta}(l),\boldsymbol{\vartheta}^s(l),\mathcal{S}(l)) \\ &\quad +\eta(l)\left(\boldsymbol{\delta M}(l)-\widehat{w(l)}\beta_h(l)\mathbf{e}^T_{t(l)}\right)+\boldsymbol{r}(l).\end{aligned} \tag{9}$$

$\boldsymbol{\vartheta}^s$ Adaptation. The random variable $s(l)$ is defined in algorithm 2 and can be split into three components $s(l) = s_{\mathcal{K}}(l) + s_1(l) + s_2(l)$. Here $s_{\mathcal{K}}(l) = s(l)\mathbb{1}(c(l) \in \mathcal{K})$ denotes the contribution of the *recurring objects*. However, the contribution of the *rare objects* is not directly $\mathcal{F}(l)$-measurable. Therefore we split it in two parts, 1) $\mathcal{F}(l)$-measurable part, $s_1(l) = \theta^s_{t(l)}(l)\mathbb{1}(c(l) \notin \mathcal{K})$, and 2)'noise', $s_2(l) = (s(l) - \theta^s_{t(l)})\mathbb{1}(c(l) \notin \mathcal{K})$. The contribution of object c to the expectation(w.r.t. $\mathcal{F}(l)$) size rate is given as,

$$\begin{aligned}g^s_c(\boldsymbol{\vartheta}(l),\boldsymbol{\vartheta}^s(l),\psi_c(l)) &= \theta^s_t(l)\mathbb{E}_l q^m_c(l)+\theta_t(l)(1-\mathbb{E}_l q^m_c(l)) \\ &\quad -\psi^0_c(l)\mathbb{E}_l q^h_c(l)-\psi^1_c(l)\mathbb{E}_l q^{hs}_c(l),\ \forall c\in\mathcal{K}.\end{aligned} \tag{10}$$

Finally, we have the conditional expectation of the update w.r.t. $\mathcal{H}(l)$ given as (full expression later),

$$\boldsymbol{g}^s\left(\boldsymbol{\vartheta}(l),\boldsymbol{\vartheta}^s(l),\mathcal{S}(l)\right) = \mathbb{E}_l\left(w(l)(s^*_{t(l)}-s_{\mathcal{K}}(l)-s_1(l))e^T_{t(l)}\right). \tag{11}$$

The adaptation of $\boldsymbol{\vartheta}^s$ for the f-TTL case (d-TTL is trivial), including similar definitions of the martingale difference, $\boldsymbol{\delta M}^s(l)$, and the reflection error, $\boldsymbol{r}^s(l)$, as

$$\begin{aligned}\boldsymbol{\vartheta}^s(l+1) =& \boldsymbol{\vartheta}^s(l)+\eta_s(l)\boldsymbol{g}^s\left(\boldsymbol{\vartheta}(l),\boldsymbol{\vartheta}^s(l),\mathcal{S}(l)\right) \\ &+\eta_s(l)\left(w(l)s_2(l)\boldsymbol{e}^T_{t(l)}+\boldsymbol{\delta M}^s(l)\right)+\boldsymbol{r}^s(l).\end{aligned} \tag{12}$$

A.2 Static Analysis

Here we analyze the caching process with fixed parameters $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^s$ which plays a crucial role in the proof of our main result. As discussed earlier fixing the latent parameters inadvertently fixes the above parameters. Using Lyapunov techniques we show that the static system converges fast towards a stationary distribution. However, using renewal theory [45] we obtain the properties of the system under stationary distribution, as it is difficult to compute the stationary distribution of the system.

Strong Mixing. The first step towards the convergence of the proposed dynamic algorithm is to show that the static system converges in an almost sure sense. We show that the stochastic process $\{\mathcal{S}(l)\}$ converges almost surely and the convergence happens at a geometric rate. The system under static setting forms a discrete time Markov chain (DTMC),

$$\boldsymbol{\mathcal{S}} = \{\mathcal{S}(l) \equiv (\boldsymbol{\Psi}(l), Z(l)) : l \in \mathbb{N}\}$$

on state space $\mathfrak{S} = [\mathbf{0},\boldsymbol{\theta}] \times [\mathbf{0},\boldsymbol{\theta}^s] \times [\mathbf{0},\boldsymbol{\theta}] \times [K+T]$.

Let the Markov kernel associated with this Harris chain be $\mathcal{P}(x, A)$, for all $A \subseteq \mathcal{B}(\mathfrak{S})$ (the Borel set on the space), $x \in \mathfrak{S}$. Further define the n-step kernel

$$\mathcal{P}^n(x, A) \equiv \mathbb{P}(\mathcal{S}(n) \in A|\mathcal{S}(0) = x).$$

Theorem 2. *Under Assumption 2 with $\|\mathbf{L}\|_\infty$-rarity condition, the Harris chain $\boldsymbol{\mathcal{S}}$ over state space $\mathfrak{S}$ admits a unique invariant measure μ^*. Furthermore, there exists constant $C > 0$ and $\gamma \in (0, 1)$ for each measurable and bounded function $f : \mathfrak{S} \to \mathbb{R}^d$, with finite d, the following holds*

$$\|\mathbb{E}f(S(n)) - \mathbb{E}_{S\sim\mu^*}f(S)\| \le C\gamma^n,\ \forall n \in \mathbb{N}. \tag{13}$$

Proof. In a Harris chain over $\mathfrak{S}$, a set $A \in \mathcal{B}(\mathfrak{S})$ is called a *small* set if there exists an integer $m > 0$ and a non zero measure ν_m such that for all $x \in A$ and for all $B \in \mathcal{B}(\mathfrak{S})$, $\mathcal{P}^m(x, B) \ge \nu_m(B)$. By showing that under the arrival process of Assumption 2, $\mathfrak{S}$ itself is a *small* set with non zero measure ν_m we can guarantee that there exist a unique invariant measure μ^* and further the chain $\boldsymbol{\mathcal{S}}$ is uniformly ergodic, i.e.

$$\|\mathcal{P}^n(S(0),\cdot) - \mu^*\|_{TV} \le (\gamma')^{n/m},\ \forall S(0) \in \mathfrak{S}.$$

Here $\gamma' = 1 - \nu_m(\mathfrak{S})$ (See Chapter 16 in [46]). Finally from the properties of the total variation norm we obtain that for any measurable and bounded (from above and below) f the above inequality (13) holds with $C = 2\sup\|f\| < \infty$ and $\gamma = (\gamma')^{1/m} \in (0, 1)$.

We show the content request process, i.e. the combination of the arrival process and the labeling process, possesses a structure that enables the wiping out of history. Our argument is based on two separate cases which are complementary to each other and results in wiping out of the effect of the initial states in two separate manner. Consider any object $c \in \mathcal{K}$. As the DTMC representing the process Z is irreducible and aperiodic, there exists some finite m_0 such that $\delta' = \min_{z\in[K+T]} \mathbb{P}(Z(_{m_0-1}) = c|Z(0) = z) > 0$.

C1 :$\mathbb{P}(X \ge L_{\max}) \ge \delta_1 > 0$.

In this case an inter arrival time of at least $L_{\max}$ effectively wipes out the history. Consider the system state

$$\mathcal{S}_{rec} = \left(\boldsymbol{\Psi} = (0, \theta^s_{c_t}, \theta_{c_t})\boldsymbol{e}_c^{|\mathcal{K}|}, Z = c\right).$$

The event, that the $(m_0 - 1)$-th inter arrival time is greater than $L_{\max}$ and the state of the labeling DTMC is c at the $(m_0 - 1)$-th step, i.e. $E = \{X(m_0 - 1) > L_{max}\} \cap \{Z(m_0 - 1) = c\}$, which happens with probability at least $\delta'_1\delta_1 > 0$, results in $\mathcal{S}(m_0 - 1) = \mathcal{S}_{rec}$. Therefore,

$$\mathcal{P}^{(m_0-1)}\left(\mathcal{S}(m_0 - 1) = \mathcal{S}_{rec}\right) \ge \delta'_1\delta_1.$$

Finally, letting $\nu_{m_0}(\cdot) = \delta'_1\delta_1\mathcal{P}(\mathcal{S}_{rec},\cdot)$, we have $\mathfrak{S}$ to be a *small* set with non zero measure ν_{m_0}. So the theorem holds for $\gamma = (1 - \delta'_1\delta_1)^{1/m_0}$.

C2 :$\mathbb{P}(X \ge L_{\max}) = 0$ and for TPM P, $\exists c, c' \in [K + T]$ s.t. $P(c, c) > 0$ and $P(c', c') > 0$.

The first part of the condition ensures that inter arrival time is bounded away from zero with positive probability, which will be used, alongside the existence of self loop, in wiping out the history. Applying Paley-Zygmund inequality, under this condition we obtain, $\mathbb{P}(X > 1/2\lambda) \ge \frac{1}{4\lambda^2 L^2_{\max}} \equiv \delta_2$. Let $\mathbb{P}(\sum_{i=1}^{n_0} X_i > L_{max}) \ge \delta_2^{n_0} > 0$ for integer $n_0 \equiv 2L_{max}\lambda$, where X_is are independent copies of the inter arrival time.

Consider two states c, c' which have self loops, i.e. $P(c, c) > 0$ and $P(c', c') > 0$. Due to irreducibility and aperiodicity of the process Z there exists a finite integer m_1 such that $\delta'_2 = \min_{z\in[K+T],d\in\{c,c'\}} \mathbb{P}(Z(m_1) = d|Z(0) = z) > 0$. Let E_1 be the event that; (i) the process Z reaches state c, and (ii) the process Z remains in the state c such that all other objects are evicted from both levels. Also, by E_2 denote the event

$$(8) = \sum\nolimits_t \left(\sum_{\substack{c:c_t=t \\ c\in\mathcal{K}}} \widehat{w_c} p\left(Z(l),c\right) \left(h_t^* - g_c(\psi_c(l), Z(l)) \right) + \widehat{\bar{w}_t}\, p(Z(l), K+t) h_t^* \right) \mathbf{e}_t^T$$

$$(11) = \sum\nolimits_t \left(\sum_{\substack{c:c_t=t \\ c\in\mathcal{U}}} w_c p\left(Z(l),c\right) \left(s_t^* - g_c^s(\boldsymbol{\vartheta}(l), \boldsymbol{\vartheta}^s(l), \psi_c(l)) \right) + \bar{w}_t\, p(Z(l), K+t) \left(s_t^* - \theta_t^s(l) \right) \right) \mathbf{e}_t^T .$$

$\widehat{\cdot}$ differentiates between byte and object hit rates.

---

that; (i) the process Z reaches from state c to state c', and (ii) the process Z remains in the state c such that all other objects are evicted from both levels. Formally, we have the events

$$E_1 = \{Z(m_1) = c | Z(0) = z\} \cap (\cap_{i=1}^{n_0} \{Z(m_1+i) = c\}) \cap \{\sum_{i=1}^{n_0} X(m_1+i) > L_{max}\}$$

$$E_2 = \{Z(2m_1+n_0) = c' | Z(m_1+n_0) = c\} \cap \left(\bigcap_{i=2m_1+n_0}^{2(m_1+n_0)} \{Z(i) = c\} \right) \cap \{ \sum_{i=2m_1+n_0}^{2(m_1+n_0)} X(i) > L_{max} \}.$$

The event $E = E_1 \cap E_2$ happens with probability at least $(\delta_2'(\delta_2 P(c,c))^{n_0})^2 > 0$. We recall that the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^s$ are fixed. Now we consider two separate conditions depending upon the pdf of inter arrival time X; letting $\delta_2'' \equiv \mathbb{P}(X > \theta^s_{c_t'})$, 1) $\delta_2'' = 1$, and 2) $\delta_2'' < 1$.

In subcase 1 ($\delta_2'' = 1$), the content c' is always evicted before it gets into the main cache $\mathcal{C}$. Therefore, following event E the system ends in the state $\mathcal{S}(2(m_1+n_0)) = \mathcal{S}_{rec}$, irrespective of the initial state. Where

$$\mathcal{S}_{rec} = \begin{cases} \left(\boldsymbol{\Psi} = (0, \theta^s_{c_t}, \theta_{c_t}) \boldsymbol{e}_c^{|\mathcal{K}|}, Z = c \right) & \text{if } c \in \mathcal{K}, \\ (\boldsymbol{\Psi} = (\mathbf{0}, \mathbf{0}, \mathbf{0}), Z = c) & \text{o/w.} \end{cases} .$$

Finally, letting $\nu_{2(m_1+n_0)}(\cdot) = \delta_2'^2 \delta_2^2 P(c,c)^{2n_0} \mathcal{P}(\mathcal{S}_{rec}, \cdot)$, we have $\mathfrak{S}$ to be a *small* set with non zero measure $\nu_{2(m_1+n_0)}$. The theorem holds for $\gamma = (1 - \delta_2'^2 \delta_2^2 P(c,c)^{2n_0})^{\frac{1}{2(m_1+n_0)}}$.

In subcase 2 ($\delta_2'' < 1$), consider additionally (to event E) the last inter arrival time is less than $\theta^s_{c_t'}$; i.e. $E' = E \cap \{X(2(m_1+n_0)) \le \theta^s_{c_t'}\}$. The event E' happens with probability at least $(1-\delta_2'')(\delta_2'(\delta_2 P(c,c))^{n_0})^2$. Finally, the event E' takes the system into the state $\mathcal{S}'(2(m_1+n_0)) = \mathcal{S}'_{rec}$ or all initial conditions. In this case we have

$$\mathcal{S}'_{rec} = \begin{cases} \left(\boldsymbol{\Psi} = (\theta_{c_t}, 0, 0) \boldsymbol{e}_c^{|\mathcal{K}|}, Z = c \right) & \text{if } c \in \mathcal{K}, \\ (\boldsymbol{\Psi} = (\mathbf{0}, \mathbf{0}, \mathbf{0}), Z = c) & \text{o/w.} \end{cases} .$$

Here, letting $\nu_{2(m_1+n_0)}(\cdot) = (1-\delta_2'')\delta_2'^2 \delta_2^2 P(c,c)^{2n_0} \mathcal{P}(\mathcal{S}'_{rec}, \cdot)$, we have $\mathfrak{S}$ to be a *small* set with non zero measure $\nu_{2(m_1+n_0)}$. The theorem holds for $\gamma = (1-(1-\delta_2'')\delta_2'^2 \delta_2^2 P(c,c)^{2n_0})^{\frac{1}{2(m_1+n_0)}}$. □

Remark 5. Condition $C1$ is true for popular inter arrival distributions such as Poisson, Weibull, Pareto and Phase-type. The consideration of these two cases let us use the Paley-Zygmund type bound in case $C2$ without using a bound on second moment of the arrival process. Further, the presence of one self loop may be used to wipe out the history of all the other states. Therefore, two self loops in the Markov chain is sufficient in wiping out the entire history. In fact, it is necessary under the current assumptions on the inter arrival time distribution, as shown next.

Remark 6 (Non-mixing of a two-level Cache.). Now we present an example where, for some $\boldsymbol{\theta}$, $\boldsymbol{\theta}^s$ and inter arrival distribution X, ergodicity of the Harris chain does not hold if; 1) there is at most one self loop in irreducible and aperiodic DTMC representing process Z and 2) $\mathbb{P}(X \ge L_{\max}) = 0$. Let there be only recurrent content of a single type and fix $L_{\max} = 2$, $\theta = 1.5$ and $\theta^s = 0.2$. Further, let the inter arrival time X be distributed with absolutely continuous pdf and support $[0.3, 0.4]$. Consider the labeling process, Z, given by the irreducible and aperiodic Markov chain; with three states $\{a, b, c\}$ and the transition probability given as $P(a,b) = P(a,a) = 0.5$, $P(b,a) = P(b,c) = 0.5$ and $P(c,a) = 1$. For the object a, corresponding to state a, the inter arrival time distribution X_a has support of $[0.3, 1.2]$. This implies if initially a is in higher level cache it is never evicted from the higher level cache. On the contrary, if a is initially in lower level cache, it never gets into the higher level cache.

Stationary Properties. We now characterize the size rate and hit rate for the f-TTL caching algorithm under the static setting. This also encompasses the d-TTL algorithm as it is essentially identical to f-TTL for $\boldsymbol{\theta}^s = \boldsymbol{\theta}$. The analysis uses results from renewal theory. To the best of our knowledge we extend the previous works and give explicit hit rate and size rate formulas under the Markovian dependence among object requests. Note that hit rate captures both byte hit rate and object hit rate with the specific definition of $\hat{w_c}$.

Theorem 3. *Under Assumption 2 with $\|\mathbf{L}\|_\infty$-rarity condition, the f-TTL algorithm with fixed parameters $\boldsymbol{\theta} \preceq \mathbf{L}$ and $\boldsymbol{\theta}^s \preceq \mathbf{L}$, achieves the average hit rate , $\boldsymbol{h}(\boldsymbol{\theta}, \boldsymbol{\theta}^s)$ (14) a.s. and the average size rate, $\boldsymbol{s}(\boldsymbol{\theta}, \boldsymbol{\theta}^s)$ (15) a.s.*

Proof. **Hit Rate.** In order to characterize the hit rate under the arrival process in 2, we first show that the combined hit rate and virtual hit rate contribution of *rare content* arrivals becomes zero almost surely.

Claim 1. *Under fixed TTLs, $\boldsymbol{\theta} \preceq \mathbf{L}$ and $\boldsymbol{\theta}^s \preceq \mathbf{L}$, and Assumption 2 with $\|\mathbf{L}\|_\infty$-rarity condition, the hit rate of 'rare content' asymptotically almost surely (a.a.s.) equals zero .*

Proof. We first note that on the l-th arrival we obtain a hit from a rare content of type t only if it is $L_{\max}$-rare content

$$\boldsymbol{h}(\boldsymbol{\theta},\boldsymbol{\theta}^s)=\left\{\left(\sum_{\substack{c:c_t=t\\ c\in\mathcal{K}}}\widehat{w_c}\lambda_c\left(p_c^2\left(\theta_t\right)+\left(1-p_c\left(\theta_t\right)\right)p_c\left(\theta_t^s\right)\right)\right)\Big/\sum_{\substack{c:c_t=t\\ c\in\mathcal{U}}}\widehat{w_c}\lambda_c\ ,\forall t\in[T]\right\} \tag{14}$$

$$\boldsymbol{s}(\boldsymbol{\theta},\boldsymbol{\theta}^s)=\left\{\left(\sum_{\substack{c:c_t=t\\ c\in\mathcal{K}}}w_c\lambda_c\left(p_c\left(\theta_t\right)\hat{s}_c\left(\theta_t\right)+\left(1-p_c\left(\theta_t\right)\right)\hat{s}_c\left(\theta_t^s\right)\right)+\sum_{\substack{c:c_t=t\\ c\notin\mathcal{K}}}\bar{w}_t\lambda_c\theta_t^s\right)\Big/\sum_{\substack{c:c_t=t\\ c\in\mathcal{U}}}w_c\lambda_c\ ,\forall t\in[T]\right\} \tag{15}$$

Here $\hat{s}_c(\theta)=\int_0^\theta x dp_c(x)+\theta(1-p_c(\theta))$ and $\widehat{\cdot}$ differentiates between byte and object hit rates.

---

of some type t, i.e. the arrival satisfies the two conditions: 1) the l-th arrival is of a rare content c of type t and 2) the previous arrival of the same rare content c happened (strictly) greater than L units of time earlier. Therefore, the hit event from rare content on l-th arrival is upper bounded by the indicator $\sum_{t\in[T]}\beta_t(l;L_{\max})$ ($\beta_t(l)$ for brevity) and the asymptotic hit rate of all the rare objects combined is upper bounded as

$$\begin{aligned}&\lim_{n\to\infty}\frac{1}{n}\sum_{l=1}^{n}\sum_{t\in[T]}\beta_t(l)\\&=\sum_{t\in[T]}\lim_{m\to\infty}\lim_{n\to\infty}\left(\frac{1}{n}\sum_{l=1}^{m-1}\beta_t(l)+\frac{1}{n}\sum_{l=m}^{n}\beta_t(l)\right)\\&\leq\sum_{t\in[T]}\lim_{m\to\infty}\lim_{n\to\infty}\frac{1}{n-m}\sum_{l=m}^{n}\beta_t(l)=0.\text{ w.p.1}\end{aligned}$$

The first term in the addition goes to zero by first taking the limit over n for finite m, whereas, the second term goes to zero, for each type separately, as a consequence of the $\|\mathbf{L}\|_\infty$-rarity condition (1). □

Recall the arrival process notations presented earlier. To characterize the hit rates it suffices to consider only the *recurring objects*, due to the above claim. For the j-th arrival (for $j\geq 1$) of object c of type t, the indicator to hit event in cache $\mathcal{C}$ is given as $\mathbb{1}_c(j)=\mathbb{1}\left(X_c(j)\leq\theta_t\right)\mathbb{1}\left(X_c(j-1)\leq\theta_t\right)$ and (real) hit in $\mathcal{C}_s$ is $\mathbb{1}_c^s(j)=\mathbb{1}\left(X_c(j)\leq\theta_t^s\right)\mathbb{1}\left(X_c(j-1)>\theta_t\right)$. The hit from the first arrival from c is (with slight abuse of notation) $\mathbb{1}_c(0)=\mathbb{1}\left(A_c(0)\leq\max\{\psi_c^0(0),\psi_c^1(0)\}\right)$. We can represent the hit rate under f-TTL at time $u>0$ is given as,

$$h_t(u)=\frac{\sum_{\substack{c:c_t=t\\ c\in\mathcal{K}}}\widehat{w_c}\left(\frac{\mathbb{1}_c(0)}{u}+\frac{1}{u}\sum_{j=1}^{N_c(u)-1}\left(\mathbb{1}_c(j)+\mathbb{1}_c^s(j)\right)\right)}{\sum_{c:c\in\mathcal{U},c_t=t}\frac{\widehat{w_c}N_c(u)}{u}}. \tag{16}$$

Due to the strong Markov property of the regeneration cycles given by the hitting times of state c, $X_c(j)$ are i.i.d. for all $j\geq 1$ and for all object $c\in\mathcal{K}$. Therefore, it forms a renewal process [45]. As $u\to\infty$, $N_c(u)/u\to\lambda_c$ from elementary renewal theorem and from the renewal reward theorem the asymptotic hit rate for type t with f-TTL is almost surely same as given in the theorem (the expression is in a placeholder).

Size Rate. Next we try to characterize the size rate of the process, that is the size per arrival in the system. The proof follows a similar idea but the only difference being the *rare objects* have non negligible contribution towards size rate for nonzero $\boldsymbol{\theta}^s$. For the j-th arrival of the object t the time it is in the cache is given as $s_c(j)=\min\{\theta_t,X_c(j+1)\}\mathbb{1}(X_c(j)\leq\theta_t)+\min\{\theta_t^s,X_c(j+1)\}\mathbb{1}(X_c(j)>\theta_t)$. We can represent the size rate under f-TTL at time $u>0$ is given as,

$$s_t(u)=\frac{\sum_{\substack{c:c_t=t\\ c\in\mathcal{U}}}w_c\left(\frac{s_c(0)+s_c(N_c(u)-1)}{u}+\frac{1}{u}\sum_{j=1}^{N_c(u)-2}s_c(j)\right)}{\sum_{c:c\in\mathcal{U},c_t=t}\frac{w_cN_c(u)}{u}}. \tag{17}$$

However if c is a *rare content* of type t we additionally have that the first term is almost surely 0 and the second term is almost surely θ_t^s. This is because the hit rate and virtual hit rate both are equal to zero, almost surely for a rare content. This simplifies the size rate contribution of the rare object of type t to $\bar{w}_t\theta_t^s$. For the stationary objects using renewal reward theorem we can obtain the rest of the terms in the expression as given in the theorem statement (the expression is in a placeholder).

Note: At each arrival of object c we add to $s(l)$ the maximum time this object may stay in the cache while we subtract the time remaining in the timer. This subtraction compensates for the possible overshoot in the previous arrival. Therefore, up to time u the sample-path wise addition of the term $s(l)$ for a object c is identical to the total bytes-sec the object has been in the cache with a small overshoot for the last arrival. Therefore, taking weighted average w.r.t. the arrival rates for object c we obtain the size rate of the system. Also in the limit $u\to\infty$ the contribution of the overshoot towards size rate becomes zero. Therefore, the average over l of the term $s(l)$ in the equation (4) represents the size rate almost surely. □

The '$\mathbf{L}$-feasible' hit rate region for d-TTL caching and f-TTL caching under arrival 2 is given in the following corollary which follows directly from the definitions.

Corollary 2. *Under Assumption 2 with $\|\mathbf{L}\|_\infty$-rarity condition, the set of '$\mathbf{L}$-feasible' hit rate for d-TTL caching algorithm is*

$$\left\{\boldsymbol{h}(\boldsymbol{\theta},\boldsymbol{\theta}):\boldsymbol{\theta}\preccurlyeq\mathbf{L}\right\}.$$

Moreover, for f-TTL caching algorithm the set of '$\mathbf{L}$-feasible' hit rate , size rate tuples is

$$\left\{(\boldsymbol{h}(\boldsymbol{\theta},\boldsymbol{\theta}^s),\boldsymbol{s}(\boldsymbol{\theta},\boldsymbol{\theta}^s)):\boldsymbol{\theta}\preccurlyeq\mathbf{L},\boldsymbol{\theta}^s\preccurlyeq\boldsymbol{\theta}\right\}.$$

We end this section with the following proposition that gives some properties of the functions $\boldsymbol{h}(\boldsymbol{\theta}, \boldsymbol{\theta^s})$ and $\boldsymbol{s}(\boldsymbol{\theta}, \boldsymbol{\theta}^s)$.

Proposition 1. *Under Assumption 2 for each t, the functions $h_t(\theta_t, \theta_t^s)$ and $s_t(\theta_t, \theta_t^s)$ are continuously differentiable for all $\theta_t \geq 0$, $\theta_t^s \geq 0$. Moreover, the function $h_t(\theta_t, \theta_t^s)$ is 1) strictly increasing (up to value 1) in θ_t for $\theta_t \geq \theta_t^s \geq 0$ and 2) non decreasing in θ_t^s.*

Proof. The inter arrival time of any object $c \in \mathcal{K}$ is given as

$$X_c = \sum_n \mathbb{P}(\inf\{l : Z(l) = c \wedge Z(0) = c\} = n) X^{(n)},$$

where $X^{(n)}$ is the sum of n independent copies of the r.v. X (inter arrival time). The first term represents the first passage time p.m.f. for the object c under the Markovian labeling. As the p.d.f. of X is absolutely continuous by Assumption 2, for each $c \in \mathcal{K}$ the p.d.f. of X_c is absolutely continuous. Further, the terms $p_c(\cdot)$ are c.d.f. of absolutely continuous random variables (a.c.r.v.) w.r.t Lebesgue measure. This means that both $h_t(\theta_t, \theta_t^s)$ and $s_t(\theta_t, \theta_t^s)$ are both continuously differentiable for all $\theta_t \geq 0$ and $\theta_t^s \geq 0$.

The p.d.f. of X, the inter arrival time, has simply connected support by assumption. Due to the continuity property of the convolution operator we can conclude that the p.d.f. of X_c also has a simply connected support for each $c \in \mathcal{K}$. This means the c.d.f.-s $p_c(\cdot)$ are strictly increasing in θ_{c_t} till it achieves the value 1 and then it remains at 1. We have the following,

$$\begin{aligned}
&\tfrac{\partial}{\partial \theta_t}\left(p_c^2(\theta_t) + (1 - p_c(\theta_t))\, p_c(\theta_t^s)\right) \\
&\quad = (2p_c(\theta_t) - p_c(\theta_t^s))\, p_c'(\theta_t). \\
&\tfrac{\partial}{\partial \theta_t^s}\left(p_c^2(\theta_t) + (1 - p_c(\theta_t))\, p_c(\theta_t^s)\right) \\
&\quad = (1 - p_c(\theta_t))\, p_c'(\theta_t^s).
\end{aligned}$$

As $p_c'(\cdot)$ have simply connected support for all $c \in \mathcal{K}$, given $h_t < 1$ and $\theta_t \geq \theta_t^s \geq 0$ the function h_t is strictly increasing in θ_t. Also h_t is trivially non decreasing in θ_t^s. □

A.3 Dynamic Analysis

In this subsection we prove our main result which shows the d-TTL and f-TTL under arrival process oblivious adaptation achieves any '**L**-feasible' hit rate where **L** is a parametric input. Our proof relies upon the tools developed in stochastic approximation theory. See [24] and the references therein for a self-contained overview of stochastic approximation theory.

Proof of Theorem 1. We consider the d-TTL and f-TTL of Section 2.2 and the arrival process as given in Assumption 2 in Section 2.3. The proof of stochastic approximation is based on the convergence of two timescale separated system using ODE methods. It is adapted from [25], [24] and [47]. First we state the required conditions for the convergence of stochastic approximation. The conditions are standard and they ensure that 1) the expected (w.r.t. $\mathcal{F}(l)$) updates are bounded and smooth, 2) they show strong ergodicity, 4) noise effects are negligible over large time windows, 3) there is time scale separation and 4) the faster timescale is approximated well by a projected ODE with 'nice' properties. These conditions lead to the approximation of the whole system by a 'nice' projected ODE evolving in the slower timescale and its convergence depicts the convergence of the system.

Condition 1 (Adapted from [25], [24] and [47])**.** Sufficient conditions for convergence:

*A*0. The step sizes satisfy
1) $\sum_l \eta(l) = \sum_l \eta_s(l) = \infty$,
2) $\sum_l \eta(l)^2 < \infty$, $\sum_l \eta_s(l)^2 < \infty$ and
3) $\lim_{l\to\infty} \eta(l) = \lim_{l\to\infty} \eta_s(l) = \lim_{l\to\infty} \frac{\eta(l)}{\eta_s(l)} = 0$.

*A*1. The update terms at each iteration satisfy
1) $sup_l \mathbb{E}\|\widehat{w(l)} Y(l) \mathbf{e}_{t(l)}^T\|_1 < \infty$ and
2) $sup_l \mathbb{E}\|w(l) s(l) \mathbf{e}_{t(l)}^T\|_1 < \infty$.

*A*2. For any sequence $\{b(l)\}$ we define $m_b(n) = \min\{k : \sum_{l=1}^{k+1} b(l) > n\}$. For $b(\cdot) \in \{\eta_s(\cdot), \eta(\cdot)\}$ the noise terms satisfy,
1) $\lim_{n\to\infty} \sum_{l=m_b(n)}^{m_b(n+1)-1} b(l)\beta_h(l) = 0$ a.s. and
2) $\lim_{n\to\infty} \sum_{l=m_b(n)}^{m_b(n+1)-1} b(l)|s_2(l)| = 0$ a.s.

*A*3. The set H is a hyper-rectangle.

*A*4. There are non negative, measurable functions $\rho_1(\cdot)$ and $\rho_1^s(\cdot)$ of $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$ such that $\|\mathbf{g}(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s, \mathcal{S})\|_1 \leq \rho_1(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s)$ and $\|\mathbf{g}^s(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s, \mathcal{S})\|_1 \leq \rho_1^s(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s)$, uniformly over all $\mathcal{S}$. Moreover the functions $\rho_1(\cdot)$ and $\rho_1^s(\cdot)$ are bounded for bounded $(\boldsymbol{\vartheta}, \boldsymbol{\vartheta})$-set.

*A*5. There are non negative, measurable and bounded functions $\rho(\cdot)$ and $\rho^s(\cdot)$ such that $\lim_{\boldsymbol{x},\boldsymbol{y}\to\mathbf{0}} \rho(\boldsymbol{x}, \boldsymbol{y}) = 0$ and $\lim_{\boldsymbol{x},\boldsymbol{y}\to\mathbf{0}} \rho^s(\boldsymbol{x}, \boldsymbol{y}) = 0$. Moreover for all $\mathcal{S}$,

$$\|\mathbf{g}(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s, \mathcal{S}) - \mathbf{g}(\boldsymbol{\vartheta}', \boldsymbol{\vartheta}'^s, \mathcal{S})\|_1 \leq \rho(\boldsymbol{\vartheta} - \boldsymbol{\vartheta}', \boldsymbol{\vartheta}^s - \boldsymbol{\vartheta}'^s).$$

Also similar inequality holds for $\boldsymbol{g}^s(\cdot)$ with $\rho^s(\cdot)$ uniformly for all $\mathcal{S}$.

*A*6. Let $\bar{\mathbf{g}}(\cdot)$ be a continuous function of $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$. Fix any $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$. Let $\boldsymbol{\xi}(l) = \mathbf{g}(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s, \mathcal{S}(l)) - \bar{\mathbf{g}}(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s)$. For some constant $C > 0$ and all n the conditions, 1)$\mathbb{E}\boldsymbol{\xi}(n) = 0$ and 2) $\sum_{i=n}^{\infty} |\mathbb{E}_n \boldsymbol{\xi}(l)| \leq C$, hold w.p. 1.

*A*7. For each fixed $\boldsymbol{\vartheta}^s$ there exists a unique $\boldsymbol{\vartheta}(\boldsymbol{\vartheta}^s)$ which correspond to a unique globally asymptotically stable point of the projected ODE, $\dot{\boldsymbol{\vartheta}} = \bar{\mathbf{g}}(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s) + \boldsymbol{r}^s(\boldsymbol{\vartheta}, \boldsymbol{\vartheta}^s)$. Here $\boldsymbol{r}^s(\cdot)$ represents the error due to projection on H.

*A*8. The functions $\mathbf{g}^s(\boldsymbol{\vartheta}(\boldsymbol{\vartheta}^s), \boldsymbol{\vartheta}^s, \mathcal{S})$ follows inequalities similar to the ones in *A*4 and *A*5 for functions $\rho_2(\cdot)$ and $\rho_3(\cdot)$, respectively.

*A*9. There exists a $\bar{\mathbf{g}}^s(\cdot)$ which is a continuous function of $\boldsymbol{\vartheta}^s$ such that statement identical to *A*6 holds for the sequence,

$$\{\mathbf{g}(\boldsymbol{\vartheta}(\boldsymbol{\vartheta}^s), \boldsymbol{\vartheta}^s, \mathcal{S}(l)) : l \in \mathbb{N}\}.$$

*A*10. There exist a twice continuously differentiable real valued function $f(\boldsymbol{\vartheta}^s)$ such that $\bar{\boldsymbol{g}}^s(\boldsymbol{\vartheta}^s) = -\nabla f(\boldsymbol{\vartheta}^s)$ for $\boldsymbol{h}^*$ which is $(1 - 2\epsilon)$**L**-feasible.

Remark 7. In A.0 part (2) is sufficient to asymptotically wipe out the contribution from the martingale difference $\boldsymbol{\delta M}$, $\boldsymbol{\delta M}^s$. In A.0 last equality of part (3) separates the two timescale making the $\boldsymbol{\theta}$ parameter vary at a faster timescale. A.1 is sufficient for uniform integrability of the updates.

Condition A.2 is sufficient for the Kushner-Clark condition to hold for the noise sequences.

We first show that all the Assumptions hold for our system.

Lemma 2. *Under TTL upper bound* $\mathbf{L}$ *and Assumption 2 with* $\|L\|_\infty$*-Rarity condition, the conditions 1.A0 to 1.A10 hold for the f-TTL and d-TTL caching.*

Proof. This is the key lemma that allows us to use the standard results in the multi-timescale stochastic approximation. The assumption that both number of recurring objects and the number of types of rare objects in the system are finite, along with the bound $\mathbf{L}$ play a crucial role in proving most of the boundedness assumptions. Furthermore, the 'rarity' condition helps us in controlling the noise from non-stationary arrival. However, the most important part of the proof lies in the analysis of the two ODEs representing the evolution of $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$.

A0: Follows from definition of $\eta(l)$ and $\eta_s(l)$ for all $l \in \mathbb{N}$.

A1: The assumption $A1$ is satisfied because both $Y(l)$ and $s(l)$ are well behaved. Specifically,

$$\begin{aligned}\sup_l \mathbb{E}\|\widehat{w(l)}Y(l)\mathbf{e}_{t(l+1)}^T\|_1 &\le w_{\max},\\ \sup_l \mathbb{E}\|w(l)s(l)\mathbf{e}_{t(l)}^T\|_1 &\le \sup_l w(l)(\theta_{t(l)}+\psi(l))\\ &\le 2w_{\max}L_{\max}.\end{aligned}$$

A2: Recall $\beta_t(l;R)$ is the indicator of the event that: (1) the l-th arrival is of a rare object c of type t, *and* (2) the previous arrival of the same rare object c happened (strictly) less than R units of time earlier, as given in the definition of rarity condition (1). Let $R = L_{\max}$ and for brevity denote $\beta_t(l;L_{\max})$ as $\beta_t(l)$, for each type t and arrival l. The event that at l-th arrival it is a hit from a rare object, $\beta_h(l)$ is upper bounded by $\sum_{t\in[T]}\beta_t(l)$ with probability 1. Further, $s_2(l)$ is non zero only if the l-th arrival results in a virtual hit or a hit from a rare object. This implies $s_2(l)$ is also upper bounded by $2L_{\max}\sum_{t\in[T]}\beta_t(l)$ a.s. Therefore, we have

$$\begin{aligned}&\sum_{l=m_\eta(n)}^{m_\eta(n+1)-1}\eta(l)\max\{|s_2(l)|,\beta_h(l)\}\\ &\qquad\le \max\{1,2L_{\max}\}\sum_{t\in[T]}\sum_{l=m_\eta(n)}^{m_\eta(n+1)-1}\eta(l)\beta_t(l).\end{aligned}$$

From the definitions, for $\alpha\in(1/2,1)$, we obtain $m_\eta(n) = (\frac{(1-\alpha)n}{\eta_0})^{1/(1-\alpha)} \pm \Theta(1)$. Therefore, there are $O(n^{\frac{\alpha}{1-\alpha}}) = O(m_\eta(n)^\alpha)$ terms in the above summation and we have the following bound

$$\sum_{t\in[T]}\sum_{l=m_\eta(n)}^{m_\eta(n+1)-1}\eta(l)\beta_t(l) \le \sum_{t\in[T]}\frac{\eta_0}{m_\eta(n)^\alpha}\sum_{l=m_\eta(n)}^{\substack{m_\eta(n)+\\ O(m_\eta(n)^\alpha)}}\beta_t(l).$$

As n goes to ∞, $m\equiv m_\eta(n)\to\infty$ and $N_m^t \equiv m_\eta(n)^\alpha = \omega(\sqrt{m})$ for $\alpha\in(1/2,1]$ and $t\in[T]$. Therefore, due to $\|\mathbf{L}\|_\infty$-rarity condition (1) and the finiteness of T, the above term goes to 0 almost surely as $n\to\infty$. This proves the assumption with step size $\eta(\cdot)$. The assumption with step size $\eta_s(\cdot)$ can be proved similarly.

A3: Follows from definition of the set H.

A4: From the definitions $\|\mathbf{g}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s,\mathcal{S})\|_1 \le 2Tw_{\max}$ easily follows. Further $\|\mathbf{g}^s(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s,\mathcal{S})\|_1 \le 2Tw_{\max}L_{\max}$.

A5: Using equations (7) and (10) we bound the terms in assumption $A5$. Fix any state $\mathcal{S}$ and let $Z=c$ for this state. First note that given the state $\mathcal{S}$ the function $\mathbf{g}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s,\mathcal{S})$ is independent of $\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s$ so $A5$ holds for $\rho=0$.

For a given $\mathcal{S}$, the function $\mathbf{g}^s(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s,\mathcal{S})$ is linear w.r.t. $\boldsymbol{\theta}^s$ and $\boldsymbol{\theta}$. However, the term $\boldsymbol{\theta}^{\boldsymbol{s}}$ has nonlinear dependence on $\boldsymbol{\vartheta}^s$ and $\boldsymbol{\vartheta}$, with bounded slope 1 and $1/\epsilon$, respectively. It is easy to verify that the following function will satisfy the conditions in $A5$

$$\rho^s(\boldsymbol{x},\boldsymbol{y}) = w_{\max}L_{\max}\left((1+1/\epsilon)\|\boldsymbol{x}\|_1 + 2\|\boldsymbol{y}\|_1\right).$$

A6: To show the validity of Assumption $A6$ we use Theorem 2. Note that $\mathbf{g}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s,\mathcal{S})$ is a measurable and bounded function satisfying the inequality (13) for $C = 2Tw_{\max}$, w.r.t. $\|\cdot\|_1$ norm. Further, due to the existence of μ^* and the above boundedness we obtain the average as $\bar{\mathbf{g}}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s) = \mathbb{E}\mathbf{g}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s,\mathcal{S})$. From the ergodicity of the system we obtain $\|\mathbb{E}\boldsymbol{\xi}(l)\|_1 = 0$. Moreover due to uniform ergodicity we have $\|\mathbb{E}_n\boldsymbol{\xi}(l)\|_1 \le 2T\gamma^{(l-n)}$ for all $l\ge n$. The bound $\sum_{i=n}^\infty|\mathbb{E}_n\boldsymbol{\xi}(l)| \le 2T/(1-\gamma)$ follows easily.

We now show that the function $\bar{\mathbf{g}}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s)$ is a continuous function of $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$. Let $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^s$ be the corresponding parameters. We claim that,

$$\bar{\mathbf{g}}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s) = \sum_t\left(\sum_{c:c_t=t}\widehat{w_c}\pi(c)\right)(h_t^* - h_t(\theta_t,\theta_t^s))\,\mathbf{e}_t^T.$$

Here $\boldsymbol{h}(\boldsymbol{\theta},\boldsymbol{\theta}^s)$ is as given in (14). To see the validity of the claim notice that the term $\mathbb{E}\,\widehat{w(l)}Y(l)$ almost surely converges to the average hit rate of the system which is obtained in Theorem 3. From Proposition 1, for each t, $h_t(\theta_t,\theta_t^s)$ is continuous with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^s$. But we know that $\boldsymbol{\theta}$ is identical to $\boldsymbol{\vartheta}$ up to scale and $\boldsymbol{\theta}^s$ is a continuous function of both $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$ from definition. Additionally we know that composition of continuous functions are continuous. Consequently $\bar{\mathbf{g}}(\boldsymbol{\vartheta},\boldsymbol{\vartheta}^s)$ is continuous in $\boldsymbol{\vartheta}$ and $\boldsymbol{\vartheta}^s$.

A7: For the rest of the proof of assumption $A7$, for notational convenience we represent the average hit rate as $\bar{\mathbf{g}}(\boldsymbol{\vartheta})$, by fixing and dropping the term $\boldsymbol{\vartheta}^{\boldsymbol{s}}$.

The function $\bar{\mathbf{g}}(\boldsymbol{\vartheta})$ is separable in types and it is the negative derivative of a real valued continuously differentiable function $f(\boldsymbol{\theta}) = -\sum_t\int g_t(\vartheta_t)d\theta_t$. This implies that the limit points of the projected ODE, $\dot{\boldsymbol{\theta}} = \bar{\mathbf{g}}(\boldsymbol{\vartheta}) + \mathbf{r}(\boldsymbol{\theta})$ ($\boldsymbol{\vartheta}^s$ suppressed) are the stationary points $\dot{\boldsymbol{\theta}} = \mathbf{0}$, lying on the boundary or in the interior of H. We next show that the stationary point under the assumptions of 'typical' hit rate is unique.

For the fixed $\boldsymbol{\vartheta}^s$ and fixed t we can represent θ_t^s as a continuous and strictly increasing function of ϑ_t. The parameter θ_t is scaled version of ϑ_t. Further, θ_t^s is increasing in ϑ_t for $\vartheta_t^s\neq 0$ and is non decreasing in ϑ_t for $\vartheta_t^s = 0$. Therefore, due to the Proposition 1 we can conclude that each function $h_t(\theta_t,\theta_t^s)$ strictly increases (w.r.t. ϑ_t) till the function reaches value 1 and after that that it remains at 1. Intuitively, for a fixed ϑ_t^s increasing ϑ_t increases the hit rate as the objects are being cached with higher TTL values in cache $\mathcal{C}$ and cache $\mathcal{C}_s$.

Now for each type t there can be two cases, 1) we achieve target hit rate h_t^* or 2) we achieve a hit rate less than h_t^*. The achieved hit rate can not be higher than that of the target hit rate due to the monotonicity and continuity property of the functions $h_t(\theta_t, \theta_t^s)$ for a fixed $\boldsymbol{\vartheta}^s$. In case (1) we find a $\vartheta_t^* \in H$, such that $h_t(\theta_t^*, \theta_t^s) = h_t^*$. Finally, due to the assumption that $\mathbf{h}^*$ is a 'typical' hit rate, i.e. $h_t^* < 1$ for all t, and the monotonicity property of $h_t(\cdot)$ w.r.t. ϑ_t we conclude that there is a unique ϑ_t^* that achieves this $\mathbf{h}^*$. However, in case (2) the ODE gets pulled to the boundary of the set by the reflection term, i.e. $\vartheta_t = 1$ and $\theta_t = L_t$. Therefore, the parameter ϑ_t^* can be expressed as a function of ϑ_t^s.

The function $\boldsymbol{\vartheta}(\boldsymbol{\vartheta}^s)$. We now present and analyze the properties of the function $\boldsymbol{\vartheta}(\boldsymbol{\vartheta}^s)$. For each t and a target hit rate $\boldsymbol{h}^*$, the function is defined separately and implicitly as

$$\vartheta_t(\vartheta_t^s) = \min\left(\{\vartheta_t : h_t\left(L_t\vartheta_t, L_t\vartheta_t\Gamma(_{\vartheta_t,\vartheta_t^s;\epsilon})\right) = h_t^*\} \cup L_t\right).$$

Intuitively, changing the value of ϑ_t will change the hit rate from the desired value of h_t^*. Taking derivative of $h_t\left(L_t\vartheta_t, L_t\vartheta_t\Gamma(\vartheta_t, \vartheta_t^s)\right)$ w.r.t. ϑ_t, yields the bound

$$\tfrac{\partial h_t}{\partial \vartheta_t} > \sum_{\substack{c:c\in\mathcal{K}\\ c_t=t}} \widehat{w_c}\lambda_c p_c(L_t\vartheta_t)p_c'(L_t\vartheta_t) \Big/ \sum_{c:c_t=t} \widehat{w_c}\lambda_c.$$

Proposition 1 states that $h_t(\cdot)$ is continuously differentiable in θ_t and θ_t^s. By definition θ_t and θ_t^s are continuously differentiable in ϑ_t and ϑ_t^s. However, the space of continuously differentiable functions is closed under composition. Therefore, $h_t\left(L_t\vartheta_t, L_t\vartheta_t\Gamma(\vartheta_t, \vartheta_t^s)\right)$ is continuously differentiable function of ϑ_t and ϑ_t^s. We also note that 1) the p.d.f. of the inter arrival time for each object c has simply connected support and 3) for a 'typical' target $\boldsymbol{h}^*$, all t and all $\vartheta_t(\vartheta_t^s)$, we have $\sum_{c:c\in\mathcal{K},c_t=t} \widehat{w_c}\lambda_c p_c(L_t\vartheta_t)p_c'(L_t\vartheta_t) > 0$. Applying a version of global implicit function theorem, as stated in Corollary 1 in [48], we conclude that for each t there exists a unique continuously differentiable function $\tilde{\vartheta}_t(\vartheta_t^s)$ which gives $h_t\left(L_t\tilde{\vartheta}_t(\vartheta_t^s), L_t\tilde{\vartheta}_t(\vartheta_t^s)\Gamma(\tilde{\vartheta}_t(\vartheta_t^s), \vartheta_t^s)\right) = h_t^*$.[7] Finally, the pointwise maximum of a constant function and a continuous function is continuous and uniquely defined. Therefore, we conclude that for each t, $\vartheta_t(\vartheta_t^s) = \max\{\tilde{\vartheta}_t(\vartheta_t^s), 1\}$ is unique and continuous.

Special case of d-TTL. Until this point it has not been necessary to differentiate the d-TTL cache from f-TTL cache. However, any hit rate that is $\boldsymbol{L}$-feasible will be achieved under the d-TTL scheme. This is not true for the f-TTL in general. Specifically, we note that for a $\boldsymbol{L}$-feasible hit rate target the case (2) above will not happen under d-TTL caching for any type t. By definition if a hit rate is not achievable under d-TTL with $\theta_t = L_t$, for at east one t, then it is not $\boldsymbol{L}$-feasible.

f-TTL Continuation: The remaining part of this proof deals with assumptions on the f-TTL caching algorithm as for d-TTL algorithm $\boldsymbol{\vartheta}^s = \mathbf{1}$ and no further analysis is necessary.

A8: We have $\|\mathbf{g}^s(\boldsymbol{\vartheta}(\boldsymbol{\vartheta}^s), \boldsymbol{\vartheta}^s, \mathcal{S})\|_1 \leq 2Tw_{\max}L_{\max}$ directly from $A4$. The function $\boldsymbol{\vartheta}(\boldsymbol{\vartheta}^s)$ as noted earlier is component wise separable and it is continuous in each component ϑ_t^s. Therefore, using the bound in $A5$ we can prove similar bounds for $\mathbf{g}^s(\cdot)$.

A9: For the function $\bar{\mathbf{g}}^s(\boldsymbol{\vartheta}^s) = \mathbb{E}\mathbf{g}^s(\boldsymbol{\theta}(\boldsymbol{\vartheta}^s), \boldsymbol{\vartheta}^s, \mathcal{S})$ the statements identical to $A6$ hold with a different constant $TL_{\max}$. The continuity of $\bar{\mathbf{g}}^s(\boldsymbol{\vartheta}^s)$ follows from the fact that of $\|\mathbf{g}^s(\boldsymbol{\theta}(\boldsymbol{\vartheta}^s), \boldsymbol{\vartheta}^s, \mathcal{S}) - \mathbf{g}^s(\boldsymbol{\theta}(\boldsymbol{\vartheta}'^s), \boldsymbol{\vartheta}'^s, \mathcal{S})\|_1 \to 0$ as $\boldsymbol{\vartheta}^s \to \boldsymbol{\vartheta}'^s$, uniformly over $\mathcal{S}$.

Similar to the hit rate calculation we can compute the function $\bar{\mathbf{g}}^s(\boldsymbol{\vartheta}^s)$ using $\boldsymbol{s}(\boldsymbol{\theta}, \boldsymbol{\theta}^s)$ from equation (15) in Theorem 1 as follows,

$$\bar{\mathbf{g}}^s(\boldsymbol{\vartheta}^s) = \sum_t \left(\sum_{c:c_t=t}\right) w_c\pi(c)\left(s_t^* - s_t(\theta_t(\vartheta_t^s), \theta_t^s(\vartheta_t^s))\right)\mathbf{e}_t^T.$$

A10: We now analyze the projected ODE evolving in the slower timescale given by $\dot{\boldsymbol{\vartheta}}^s = \bar{\boldsymbol{g}}^s(\boldsymbol{\vartheta}^s) + \boldsymbol{r}^s(\boldsymbol{\vartheta}^s)$, where $\boldsymbol{r}^s(\boldsymbol{\vartheta}^s)$ is the reflection term. The first consequence of the separability of the function $\bar{\boldsymbol{g}}^s(\cdot)$ among types is the representation of the same function as a negative gradient of real valued continuously differentiable function $f^s(\boldsymbol{\vartheta}^s) = -\sum_t \int \bar{g}_t^s(\vartheta_t^s)d(\vartheta_t^s)$. Which readily makes the limit point of the ODE as the stationary points, i.e. the points where $\dot{\boldsymbol{\vartheta}}^s = \mathbf{0}$.

Claim 2. *Each stationary point of the slower ODE, given a $((1-2\epsilon)\boldsymbol{L})$-feasible tuple $(\boldsymbol{h}^*, \boldsymbol{s}^*)$, corresponds to a TTL pair $\left(\hat{\boldsymbol{\vartheta}}, \hat{\boldsymbol{\vartheta}}^s\right)$ and an achieved hit rate, size rate tuple $\left(\hat{\boldsymbol{h}}, \hat{\boldsymbol{s}}\right)$, that, for each type t, satisfy one of the following three conditions:*

1. *Hit rate $\hat{h}_t = h_t^*$ and size rate $\hat{s}_t = s_t^*$.*
2. *Hit rate $\hat{h}_t = h_t^*$ and size rate $\hat{s}_t > s_t^*$, while $\hat{\vartheta}_t^s = 0$.*
3. *Hit rate $\hat{h}_t = h_t^*$ and size rate $\hat{s}_t < s_t^*$, while $\hat{\vartheta}_t^s = 1$.*

Proof. The separability of the ODE in each type t enables us to analyze each coordinate of the ODE separately. First of all it is not hard to check that the points mentioned above gives $\bar{\boldsymbol{g}}^s(\boldsymbol{\vartheta}^s) = 0$, i.e. they are stationary. The first case holds as the expectation of the update becomes zero. However, for the other two cases the projection term induces stationary points on the boundary. The harder direction is to rule out existence of other stationary points under $((1-2\epsilon)\boldsymbol{L})$-feasibility of the tuple.

We now recall the properties of the convergence of the projected ODE in the faster timescale given a fixed $\boldsymbol{\vartheta}^s$. Specifically, for a fixed $\boldsymbol{\vartheta}^s$ for each t there can be two cases (1) for some unique $\boldsymbol{\theta}$ the hit rate $\boldsymbol{h}^*$ is achieved or (2) the hit rate is not achieved for types $t \in T' \subseteq T$ and $\theta_t = L_t$.

However, the scenario (2) is not possible for $(1-2\epsilon)\boldsymbol{L}$-feasible hit rates. For the sake of contradiction let us assume that there is a stationary point where $\theta_t = L_t$. This implies that $\theta_t^s = L_t$ and the cache is working in the d-TTL mode. The hit rate of object t in f-TTL is always smaller than that of d-TTL with same θ_t value. This means the 'typical' hit rate $\boldsymbol{h}^*$ can not be $(1-2\epsilon)\boldsymbol{L}$-feasible as it lies on the boundary of $\boldsymbol{L}$-feasible region.

We are now left with the scenario that the hit rate is achieved, i.e. $\hat{h}_t = h_t^*$. We can further differentiate three cases (1a) $\theta_t \leq (1-3\epsilon/2)L_t$, (1b) $(1-3\epsilon/2)L_t < \theta_t < (1-\epsilon/2)L_t$ and (1c) $\theta_t \geq (1-\epsilon/2)L_t$. First of all we can rule out (1c) from the same logic as scenario (2). In case (1a) due to $\boldsymbol{L}$-feasibility we achieve the desired tuple as there exist some stationary point in our mentioned space. In the case (1b) external force works on the ODE where we start

[7] In corollary 1 [48], the condition (1) is true due to the existence of a ϑ_t for all ϑ_t^s for 'typical' hit rate. The condition (3) holds as the functions $h_t(\cdot)$ are all measurable.

moving f-TTL towards d-TTL. Nonetheless, the projected ODE for $\boldsymbol{\vartheta}^s$ tracks the average size of the system and reaches a stationary point only if the desired size rate is achieved. Consequently, it achieves the targeted tuple with $\theta_t > L_t(1-2\epsilon)$. Therefore, none of the above three cases, (1a), (1b) and (1c), occurs. □

Finally, to conclude the proof we have shown that for all t the function $\bar{g}_t^s(\vartheta_t^s)$ are continuously differentiable in ϑ_t^s. Recall that the function $\vartheta_t(\vartheta_t^s)$ is defined as $\vartheta_t(\vartheta_t^s) = \max\{\tilde{\vartheta}_t(\vartheta_t^s), 1\}$ where $\tilde{\vartheta}_t(\vartheta_t^s)$ is a continuously differentiable function of ϑ_t^s. For any $(1-2\epsilon)\boldsymbol{L}$-feasible hit rate target we know that $\vartheta_t(\vartheta_t^s) < L_t$ implying $\vartheta_t(\vartheta_t^s) \equiv \tilde{\vartheta}_t(\vartheta_t^s)$.[8] Therefore, for each t the function $\bar{g}_t^s(\vartheta_t^s)$ is continuously differentiable (the space closed under composition) and the function $f^s(\boldsymbol{\vartheta}^s)$ is twice continuously differentiable. □

We state the following theorem which is a direct adaptation from the works by [25], [47] and [24].

Theorem 4 (Adapted from [25], [47] and [24])**.** *Under the validity of A0 to A10 in Condition 1, there is a null set $\mathcal{N}$ such that for any sample path $\omega \notin \mathcal{N}$, $(\boldsymbol{\vartheta}(l), \boldsymbol{\vartheta}(l))^{\omega}$ almost surely converges to a unique compact and connected subset of the stationary points given in the Claim 2, as $l \to \infty$.*

This completes the proof by noting that the stationary points yield only the scenarios mentioned in the main theorem statement. □

Remark 8 (Theorem 4). The assumption of the target hit rate $\mathbf{h}^*$ to be 'typical' as well as the condition that the inter arrival time p.d.f. is simply connected is introduced to prove convergence of the algorithm to a unique $\boldsymbol{\theta}^*$. If we relax these assumptions the convergence for d-TTL still happens and achieves the targeted hit rate. But it converges to a set of stationary points of the projected ODE for the faster timescale. However the two timescale analysis fails to hold as the uniqueness of the stationary point in the faster projected ODE is a standard assumption in the existing literature. On the contrary the absolute continuity of the p.d.f. of the inter arrival time is essential for the dynamics of the system to be continuous. Relaxing this assumption will deem the analysis infeasible even for d-TTL. The rarity condition can not be weakened significantly without violating Kushner-Clark condition on the noise sequence. We, also, note as the target approach the higher hit rate regions the f-TTL algorithm starts becoming less aggressive in meeting the size rate. In this case f-TTL starts moving towards d-TTL algorithm by the use of threshold function $\Gamma(\cdot,\cdot;\epsilon)$.

A.4 Poisson Arrival, Independent Labeling

In this section we prove Lemma 1 which states that full filtering yields the best size rate if the hit rate is achievable in the full filtering mode. We also give a guideline about setting the parameter for the algorithm for a given hit rate target in Lemma 4.

[8]The assumption $A10$ fails to hold for the the hit rate and size rate tuples at the 'boundary' of $\boldsymbol{L}$-feasible region as $\vartheta_t(\vartheta_t^s)$ is no longer continuously differentiable. This implies that the convergence may happen to any 'asymptotically stable' point in H. These points are really hard to characterize for the complex caching process. Refer [24] for details.

Proof of Lemma 1. From the previous discussions, it suffices to prove the lemma for each type separately. So we consider the objects of type t and drop the subscript t (at times) for simplifying notations without loss of clarity. Given target hit rate h^*, let $(\hat{\theta}, \hat{\theta}^s)$ and $(\theta, 0)$ be the TTLs for the two f-TTL caching algorithms. For any TTL pair (θ, θ^s) we have the hit rate of recurring object c as

$$h_c(\theta,\theta^s) = \left(\left(1-e^{-\lambda_c\theta}\right)^2 + e^{-\lambda_c\theta}\left(1-e^{-\lambda_c\theta^s}\right)\right).$$

Further, for Assumption 3 the size rate of the object c can be expressed as $s_c(\theta,\theta^s) = \lambda_c^{-1}h_c(\theta,\theta^s)$. The combined hit rate and size rate for type t can be given as

$$h(\theta,\theta^s) = \frac{\sum_{c:c_t=t} w_c\pi_c h_c(\theta,\theta^s)}{\sum_{c:c_t=t} w_c\pi(c)},\ s(\theta,\theta^s) = \frac{\sum_{c:c_t=t} w_c\pi_c s_c(\theta,\theta^s)}{\sum_{c:c_t=t} w_c\pi(c)}.$$

The difference of size rate for a *recurring object* c among this two caches can be given as

$$\Delta_c = h_c(\hat{\theta},\hat{\theta}^s) - h_c(\theta,0).$$

Due to monotonicity of hit rate we have $\theta > \hat{\theta}$ for 'typical' target hit rate. Intuitively, as the total hit rate for the type remains fixed more filtering increase the hit rate of the popular objects in f-TTL with $(\theta,0)$ compared to the other, whereas for the unpopular objects the order is reversed. Building on the intuition, for $\delta_1 = 1 - \frac{\hat{\theta}^s}{\hat{\theta}} \in [0,1)$, $\delta_2 = \frac{\theta}{\hat{\theta}} - 1 > 0$ and $x_c = e^{-\lambda_c\hat{\theta}}$ we can represent Δ_c as $q(x_c)$, where $q(\cdot)$ is

$$q(x) = (1-x)^2 + x(1-x^{1-\delta_1}) - (1-x^{1+\delta_2})^2.$$

Analyzing this polynomial we can conclude that the polynomial has exactly one root x^* in the interval $(0,1)$. Therefore, it is easy to check that for $x_c \in [x^*,1)$, $\Delta_c \geq 0$ and for $x_c \in (0,x^*)$, $\Delta_c < 0$. From definition of x_c, we deduce that there is a c^* such that $\Delta_c < 0$ for $c \leq c^*$ and $\Delta_c \geq 0$ otherwise.

W.l.o.g. let c be ordered (otherwise relabel) in the decreasing order of popularity ($c=1$ most popular). Therefore, first due to the equality of the hit rate for the two scenario we have

$$\sum\nolimits_{c=c^*+1}^{|\mathcal{K}_t|} \hat{w}_c\pi_c\Delta_c = \sum\nolimits_{c=1}^{c^*} \hat{w}_c\pi_c|\Delta_c|, \tag{18}$$

where $\hat{w}_c = 1$ if object hit rate is considered else, in case of byte hit rate, $\hat{w}_c = w_c$. Further, for type t we have the

difference in the average size rate over all objects as,

$$\left(\sum_{c:c_t=t} w_c\lambda_c\right)\Big(s(\hat{\theta},\hat{\theta}^s)-s(\theta,0)\Big)$$

$$=\sum\nolimits_{c=1}^{|\mathcal{K}_t|} w_c\Delta_c+\bar{w}_t\alpha_t\hat{\theta}^s, \tag{19a}$$

$$=\sum\nolimits_{c=c^*+1}^{|\mathcal{K}_t|} w_c\Delta_c-\sum\nolimits_{c=1}^{c^*} w_c|\Delta_c|+\bar{w}_t\alpha_t\hat{\theta}^s, \tag{19b}$$

$$\geq \tfrac{w_{c^*+1}}{\hat{w}_{c^*+1}}\sum\nolimits_{c=c^*+1}^{|\mathcal{K}_t|}\hat{w}_c\Delta_c-\tfrac{w_{c^*}}{\hat{w}_{c^*}}\sum\nolimits_{c=1}^{c^*}\hat{w}_c\pi_c|\Delta_c|+\bar{w}_t\alpha_t\hat{\theta}^s, \tag{19c}$$

$$\geq \tfrac{w_{c^*+1}}{\hat{w}_{c^*+1}}\left(\sum\nolimits_{c=c^*+1}^{|\mathcal{K}_t|}\hat{w}_c\Delta_c-\sum\nolimits_{c=1}^{c^*}\hat{w}_c\pi_c|\Delta_c|\right)+\bar{w}_t\alpha_t\hat{\theta}^s, \tag{19d}$$

$$\geq \tfrac{w_{c^*+1}}{\hat{w}_{c^*+1}}\left(\sum\nolimits_{c=c^*+1}^{|\mathcal{K}_t|} w_c\Delta_c-\frac{1}{\pi_{c^*}}\sum\nolimits_{c=1}^{c^*} w_c\pi_c|\Delta_c|\right)+\bar{w}_t\alpha_t\hat{\theta}^s, \tag{19e}$$

$$=\tfrac{w_{c^*+1}}{\hat{w}_{c^*+1}}\sum\nolimits_{c=c^*+1}^{|\mathcal{K}_t|} w_c\left(1-\tfrac{\pi_c}{\pi_{c^*}}\right)\Delta_c+\bar{w}_t\alpha_t\theta_1^s \tag{19f}$$

$$\geq \bar{w}_t\alpha_t\hat{\theta}^s. \tag{19g}$$

The equalities (19a) and (19b) are trivial, as well as the inequalities (19c) and (19d) in case of byte hit rate. Whereas, for object hitrate, the inequalities (19c) and (19d) hold as, by assumption, the size of a recurrent object is non-decreasing w.r.t. probability π_c, i.e. $w_c \geq w_{c'}$ for $c < c'$. The inequality (19e) holds as $\frac{\pi_c}{\pi_{c^*}} \geq 1$ for $c \leq c^*$ and the equality (19f) is due to the relation (18). Finally, the last inequality (19g) follows from, $\frac{\pi_c}{\pi_{c^*}} \leq 1$ and $\Delta_c > 0$ for $c > c^*$. This completes the proof. □

The above lemma strengthens the statement in Theorem 4 under the assumption of Poisson traffic with independent labeling, as stated in Corollary 1, which we prove next.

Proof of Corollary 1. We know from Theorem 1 that for any type t, the f-TTL algorithm converges to one of the three scenarios. In scenario 1 and 3, the statement of the corollary immediately holds. We now show by contradiction that convergence to scenario 2 never happens. Suppose the algorithm converges to scenario 2. The algorithm achieves average hit rate h_t^* and average size rate $\tilde{s}_t > s_t^*$, by the supposition. However, as the tuple $(\boldsymbol{h}^*, \boldsymbol{s}^*)$ is $(1-2\epsilon)\boldsymbol{L}$-feasible only the two following cases take place.

In the first case there exists a tuple (θ_t, θ_t^s), with $\theta_t^s > 0$, such that for type t we achieve the tuple (h_t^*, s_t^*) (a.a.s.). Due to Lemma 1, this implies the size rate under scenario 2, $\tilde{s}_t \leq s_t^*$. This leads to a contradiction.

In the second case (complement of the first) there exists a tuple $(\tilde{\theta}_t, 0)$ that achieves (h_t^*, s_t^*). However, as $h_t^* < 1$ ($\boldsymbol{h}^*$ is 'typical'), under full filtering ($\theta_s^t = 0$) the hit rate monotonically increases with the TTL value. Therefore, there is a unique value of θ_t for which h_t^* is achieved (a.a.s.). Therefore, convergence in scenario 2 happens to pair $(\tilde{\theta}_t, 0)$ (a.a.s.) and the average size rate $\tilde{s}_t = s_t^*$. However this is again a contradiction. □

A.5 Poisson Arrival and Zipfian Labeling

Finally, we present a guideline for parameter setup for Poisson arrival and Zipfian distribution—a popular model in caching literature. Consider the following Zipfian popularity distribution where the object label distribution is as follows.

Assumption 4. Poisson Arrival and Zipfian Labeling:

- The inter arrival times are i.i.d. and follows exponential distribution with rate $\lambda = n\lambda_0$, for constant $\lambda_0 > 0$.
- Let $\pi(c)$ be the probability of a recurring object c. Then for each type $t \in [T]$, the marginal probability of recurring objects is $q_t = \sum_{c:c_t=t, c\in\mathcal{K}} \pi(c)$. Whereas, on each turn, some rare object of type t is chosen with probability α_t and following the condition in (1).
- Arrange the recurrent objects of type t in decreasing order of the popularity $\mathcal{K}_t = (c : c_t = t, c \in \mathcal{K})$. For each type t, the popularity of recurrent objects follows a Zipfian distribution with parameter β_t. Specifically, let $c_k = \mathcal{K}_t(k)$ be the k-th popular recurrent object of type t, then $\pi(c_k) = q_t Z_t^{-1} k^{-\beta_t}$, for $k = 1, \ldots, K_t$.
- All the objects have unit size, i.e. $w_c = 1$ for all c.

We next discuss the choice of appropriate truncation parameter $\mathbf{L}$ for Poisson arrival and Zipfian distribution.

Lemma 3. *Suppose, for some constant $\gamma > 0$, the parameter $\mathbf{L}$ satisfy,*

$$\mathbf{L} = \Big(L_t : t \in [T], L_t = \tilde{\Omega}\Big(n^{-(1-\frac{\beta_t\gamma}{\beta_t-1})}\Big)\Big).$$

Then, under Assumption 4 with $\|\mathbf{L}\|_\infty$-rarity condition, for both d-TTL and f-TTL algorithms an object hit rate $\boldsymbol{h}^* = (1-\Omega(n^{-\gamma}))\mathbf{1}^T$, is $\mathbf{L}$-feasible.

Proof of Lemma 3. Let c be a type t object and $\lambda_c = \lambda q_t Z_t^{-1} c^{-\beta_t}$. For d-TTL algorithm the hit rate for type t objects is given as,

$$h_t(\theta_t) = \frac{q_t}{\alpha_t + q_t}\left(\sum\nolimits_{c=1}^{|\mathcal{K}_t|} Z_t^{-1} c^{-\beta_t}(1-e^{-\lambda_c\theta_t})\right).$$

Further define $r_t^* = 1 - \frac{h_t^*}{1+\alpha_t/q_t}$ and $\delta r_t = \frac{r_t^*}{\log(n)}$. For $c_t^* = (\delta r_t Z_t(\beta_t - 1))^{-\frac{1}{\beta_t - 1}}$, we have $\sum_{c=c_t^*+1}^{|\mathcal{K}_t|} \lambda_c \leq \delta r_t$.

$$\sum\nolimits_{c=1}^{|\mathcal{K}_t|} Z_t^{-1} c^{-\beta_t} e^{-\lambda_c L_t} \leq e^{-\lambda_{c_t^*} L_t} + \delta r_t.$$

Choosing $L_t = \frac{1}{\lambda_{c_t^*}} \ln\left(\frac{1}{0.99 r_t^* - \delta r_t}\right)$, $e^{-\lambda_{c_t^*} L_t} = (0.99 r^* - \delta r_t)$ and the hit rate for TTL L_t, is strictly greater than h_t^*. Furthermore, as $r_t^* = \mathcal{O}(n^{-\gamma})$ and $\lambda = n\lambda_0$, we have

$$L_t = \Omega\left(n^{-(1-\frac{\beta_t\gamma}{\beta_t-1})}(\log n)^{1+\frac{\beta_t}{\beta_t-1}}\right).$$

For f-TTL we arrive at the order-wise same result for L_t by observing the hit rate expression for the f-TTL is only constant times away from that of d-TTL. □

Remark 9. Lemma 3 implies that under Poisson arrival with arrival rate $\lambda = \Omega(n)$ and Zipfian popularity distribution with $\min_t \beta_t > 1$, we can achieve arbitrarily high hit rates using parameter $\boldsymbol{L}$ of the order $O(1)$ for large enough n and appropriate rarity condition. The above lemma considers only unit sized object, but can be extended to more general object sizes that are bounded.

B. BYTE HIT RATE PERFORMANCE OF D-TTL AND F-TTL

In this section, we evaluate the byte hit rate performance and convergence of d-TTL and f-TTL using the experimental setup described in Section 6.1.

B.1 Hit rate performance of d-TTL and f-TTL

To obtain the HRC for d-TTL, we fix the target byte hit rate at 80%, 70%, 60%, 50% and 40% and measure the hit rate and cache size achieved by the algorithm. Similarly, for f-TTL, we fix the target hit rates at 80%, 70%, 60%, 50% and 40%. Further, we set the target size rate of f-TTL to 50% of the size rate of d-TTL. The HRCs for byte hit rate is shown in Figure 9.

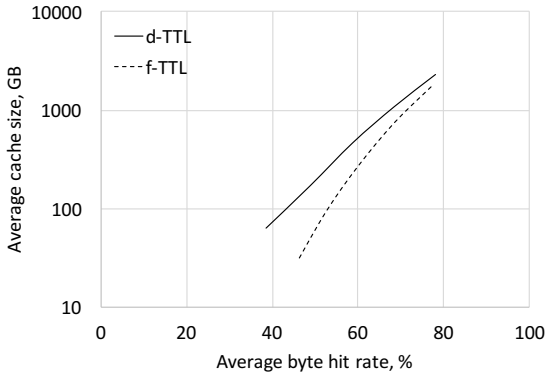

Figure 9: Hit rate curve for byte hit rates.

From 9, we see that, for a given hit rate, f-TTL requires lesser cache space than d-TTL. On average, f-TTL requires a cache that is 39% smaller than d-TTL to achieve the same byte hit rate. Further note that achieving a specific byte hit rate value requires more cache size than achieving the same value for object hit rate. For instance, the d-TTL algorithm requires a cache size of 469GB to achieve a 60% byte hit rate, whereas a 60% object hit rate is achievable with a smaller cache size of 6GB (Figure 5). This discrepancy is due to the fact that popular objects in production traces tend to be small (10's of KB) when compared to unpopular objects that tend to be larger (100's to 1000's of MB).

B.2 Convergence of d-TTL and f-TTL for byte hit rates

In this section we measure the byte hit rate convergence over time, averaged over the entire time window and averaged over 2 hour windows for both d-TTL and f-TTL. We set the target hit rate to 60% and a target size rate that is 50% of the size rate of d-TTL.

From Figures 10 and 11, we see that d-TTL has a cumulative error of less than 2.3% on average while achieving the target byte hit rate and f-TTL has a corresponding error of less than 0.3%. Moreover, we see that both d-TTL and f-TTL tend to converge to the target hit rate, which illustrates that both d-TTL and f-TTL are able to adapt well to the dynamics of the input traffic. We also observe that the average byte hit rates for both d-TTL and f-TTL have higher variability compared to object hit rates (Figures 6 and 7, due to the the fact that unpopular content

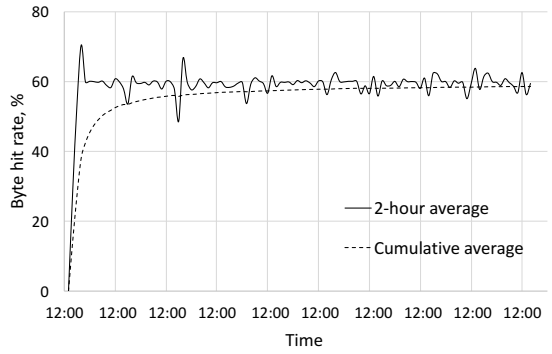

Figure 10: Byte hit rate convergence over time for d-TTL; target byte hit rate=60%.

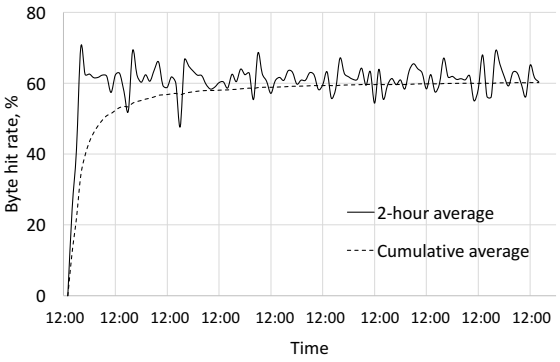

Figure 11: Byte hit rate convergence over time for f-TTL; target byte hit rate=60%.

in our traces have larger sizes, and the occurrence of non-stationary traffic can cause high variability in the dynamics of the algorithm.

In general, we also see that d-TTL has lower variability for both object hit rate and byte hit rate compared to f-TTL due to the fact that d-TTL does not have any bound on the size rate while achieving the target hit rate, while f-TTL is constantly filtering out non-stationary objects to meet the size rate target while also achieving the target hit rate.